Redefining How Pharmaceutical Innovation Gets Done

by

Vahid Montazerhodjat
S.M., Electrical Engineering and Computer Science, MIT, 2013

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in Electrical Engineering and Computer Science
at the Massachusetts Institute of Technology

February 2016

Signature of Author: Signature redacted

Vahid Montazerhodjat
Department of Electrical Engineering and Computer Science
December 29, 2015

Certified by: Signature redacted

Andrew W. Lo
Charles E. and Susan T. Harris Professor
Director, Laboratory for Financial Engineering
Thesis Supervisor

Accepted by: Signature redacted

Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Committee for Graduate Students

Redefining How Pharmaceutical Innovation Gets Done

by Vahid Montazerhodjat

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in Electrical Engineering and Computer Science

Abstract

The productivity of research and development in the bio-pharmaceutical industry has been constantly declining since the early 2000's. One possible reason is that biomedical projects are risky, take a long time, and require significant investment. Hence, substantial capital has been shifted away from the bio-pharmaceutical industry to other industries that are perceived less risky, creating a funding gap for early-stage pharmaceutical R&D. Here, we investigate and improve upon a novel financing technique that has been proposed to facilitate the R&D funding in the bio-pharmaceutical industry. This new financing method is a clear example of rapidly evolving innovation in the financial industry, from which the bio-pharmaceutical industry can benefit tremendously.

Apart from funding challenges, pharmaceutical companies have to clear regulatory hurdles before they can commercialize their treatments. These drug-regulatory standards require a specific balance of benefits vs. risks for a therapy to be approved, and do not currently take into account the severity of the disease that the therapy is targeting. In the second part of this thesis, we propose an objective and quantitative Bayesian decision analysis framework to incorporate patients' feedback into the drug-approval process, and propose adjustment to the approval standards based on disease severity.

When launching their drug, pharmaceutical companies set the drug's price such that expected revenues offset the costs of all the projects, failed or successful, that were pursued in order to lead to this successful treatment resulting in costly treatment. Recently, some highly curative therapies with high price tags have emerged for diseases with large prevalence, such as hepatitis C. These high prices, coupled with the large size of the patient population, have created an unsupportable financial burden for insurance companies in order to cover the broadest patient population who could benefit from these drugs. Despite delivering breakthrough discoveries, the pharmaceutical companies producing these drugs have experienced a public backlash due to drug prices. In the last part of this dissertation, we introduce a new financing paradigm to address the issue of high aggregate costs for these highly curative therapies.

Thesis Supervisor: Andrew W. Lo
Title: Charles E. and Susan T. Harris Professor
Director, Laboratory for Financial Engineering

Acknowledgments

First and foremost, I am grateful to my adviser, Andrew Lo, for his endless support and guidance since the first day we started working together. Working with Andrew has been an amazing journey, filled with intellectual joy and curiosity, and never with a dull moment. Andrew has been and will always be a role model in my career, and I'm excited to learn much more from him in the coming years. I'd also like to thank my doctoral committee members, John Frishkopf and Devavrat Shah, who gave me invaluable feedback for polishing this work. I'd like to give a shout-out to John who has been my mentor since last year, and has been quite patient with my sometimes slow learning. John, you are an amazing person and a great friend. I greatly appreciate it that you always make some time to discuss different topics and to answer questions despite your extremely busy work schedule. The fact that you are so down to earth makes working with you even more pleasant. I have looked up and will definitely look up to you in my professional and personal lives, and hope that we continue our collaboration going forward. All the remaining errors in this dissertation are solely mine and do not necessarily represent the views and opinions of the individuals acknowledged here or any of their institutions.

I'm thankful for all the kind help and assistance that the LFE's amazing administrative assistants have continually provided us with. Thank you Allie, Jayna, and Patsy. Our projects would have gone nowhere had it not been for you and your great help.

Last but not least, I'd like to express my gratitude to my parents, sister, brother, and my best friend, Alexa, who have unconditionally supported me in finishing this chapter of my studies. I am so lucky to have all of you in my life and am grateful for all your endless support and love. I owe everything to my parents, who are the most precious persons of my life, and I'd like to dedicate this dissertation to my hero—my mom—for her endless sacrifices to raise all three of us.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

DESPITE many breakthroughs in the pharmaceutical and biotechnology industries, including human genome sequencing, use of biomarkers to measure the response to drugs, invention of better diagnostic techniques, and discovery of new biologics, the productivity of the pharmaceutical industry—as measured by the number of new drugs approved per R&D spending—has constantly been declining since the early 2000's [1–4]. The scientific breakthroughs achieved over the past decade have opened many avenues to pursue before a safe and effective drug for a disease is typically discovered. The amount of information that academic researchers and pharmaceutical practitioners have obtained through the use of genetic sequencing has ballooned over the past decade because the cost of these technologies has constantly been declining [5]. However, there has not yet been a proper infrastructure to use the obtained information efficiently. The sheer number of possible paths to take, in addition to the ever-growing regulatory safety requirements, has consequently made the drug development process costlier, more challenging, riskier, and even longer [6].

Furthermore, the very scientific progress that has answered so many crucial questions about the biology of different diseases has also pushed therapeutic solutions toward more personalized medicine and has consequently created a smaller target population size for a typical marketed drug [7]. For example, the term *breast cancer* no longer refers to a single disease; it instead serves as an umbrella for a set of diseases with quite heterogeneous genetics and biological profiles, which in turn require different treatments and have different afflicted populations [8–10].[1] Therefore, even after going through a long, risky, and costly process, the revenue generated from the sale of a marketed drug may not be as large as it was a decade ago—in the case of the so-called *blockbusters*—unless pharmaceutical companies charge dramatically higher prices for these personalized therapeutics.

[1] See `http://www.nationalbreastcancer.org/types-of-breast-cancer`.

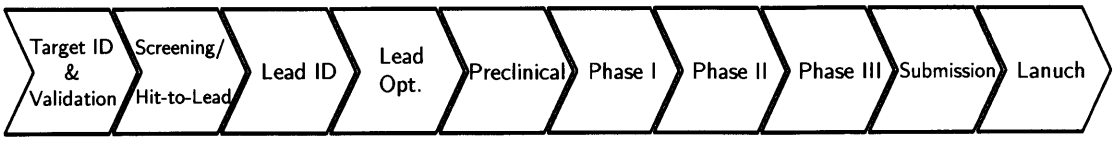

Figure 1.1. A schematic for drug research and development chain.

Note. ID: identification, Opt: optimization. The submission stage is sometimes referred to as the registration stage as well.

These factors, among many others, have led to unsatisfactory returns on the investment in early-stage drug research and discovery (R&D) made by conventional sources of financing, namely, private and public equity. Recent economic conditions, e.g., the financial crisis of 2008, and heated political debates on the fair valuation of breakthrough curative therapies might further have discouraged investors from investing in the bio-pharmaceutical industry. As a result of this growing uncertainty in the bio-pharmaceutical industry over the past decade, venture capital firms, one of the main funding sources for biotechnology start-up companies, have reduced their investment in early-stage drug R&D and increased investment in other sectors that are perceived as less risky, i.e., software and information technology (IT) [11]. For instance, the life-sciences share of venture capital investment dollars in 2013 experienced its lowest level since 2001 at 23.6% of the total dollars as documented in the National Venture Capital Association Yearbook 2014 [11]. This lowest level of investment in the life-sciences sector is accompanied by venture capital investment in the software sector receiving the highest percentage of total dollars in 2013 [11]. This outflow of capital from the life-sciences sector, in turn, has caused a severe gap in funding between early-stage drug research, usually funded through research grants provided by government agencies, and late-stage clinical development, funded by large biotechnology and pharmaceutical companies. This funding gap is usually referred to as the *valley of death*, and consists of the discovery stage up to the preclinical stage in Figure 1.1.

Fortunately, the existence of risk and uncertainty in projects is not unique to the bio-pharmaceutical industry and the notion of risk mitigation and management is at the heart of financial economics. In fact, not only is there extensive literature on risk mitigation in finance dating back decades, but there is also a well-established apparatus in the financial markets for efficient implementation of these risk-sharing methodologies. Due to the uniquely high competition in the financial industry, some of the tools developed for risk management in financial markets are extremely efficient.

Applying efficient tools borrowed from the financial industry to risk management in the pharmaceutical industry is not a new concept. In fact, some of the big pharmaceutical companies have been applying different financial risk management tools for several decades [12]. However, the extent to which financial innovation has so far helped the bio-pharmaceutical industry, especially the biotechnology companies, seems far from satisfactory compared to the capacity of these innovations. Hence, the topics and tools covered in this dissertation are unique in that they have not yet been employed in large scale in the bio-pharmaceutical industry, and that they have the potential to change the way that innovations get done in the pharmaceutical industry.

Outline. In Chapter 2, we improve upon a recently proposed method for financing drug research and discovery, specifically, the early-stage drug R&D, covering different parts of the drug discovery and development up to Phase II and Phase III in Figure 1.1 [13]. The proposed method in Chapter 2 allows bio-pharmaceutical companies to tap into the global capital markets so as to lock in funding for their innovative, yet risky projects, and to manage their exposure to different clinical compounds in their pipelines as desired. In Chapter 3, we introduce a Bayesian decision analysis framework to improve the efficiency of the drug-approval process that takes place at the Phase III and submission stages of Figure 1.1. In that chapter, we aim to design a quantitative drug-regulatory framework, in which the required safety/efficacy criteria in the drug-approval process get adjusted based on the severity of the disease that the drug under test intends to treat. Then, in Chapter 4, we propose a new method for the reimbursement of highly curative therapies to address the issue of large aggregate costs for payers and insurance companies, caused by the high price tag of these new treatments. In Figure 1.1, this problem corresponds to the launch stage, where pharmaceutical companies get reimbursed for their marketed therapeutics. Finally, we conclude the dissertation in Chapter 5.

Chapter 2

Financing Drug Research and Development

Recently, Fernandez, Stein, and Lo [13] proposed a new financing vehicle, termed megafund, for early-stage drug research and discovery. Fagnan *et al.* [14] applied the same methodology to a portfolio of drug candidates for orphan/rare diseases. We extend the megafund concept for funding drug discovery to allow for dynamic leverage in which the portfolio of candidate therapeutic assets is financed initially by mostly equity, and debt is introduced gradually as assets mature and begin generating cash flows. Leverage is adjusted so as to maintain an approximately constant level of default risk throughout the life of the fund. Numerical simulations show that applying dynamic leverage to a small portfolio of orphan drug candidates can boost the return on equity almost twofold compared to securitization with a static capital structure. Dynamic leverage can also add significant value to comparable all-equity-financed portfolios, enhancing the return on equity without jeopardizing debt performance or increasing risk to equity investors.

■ 2.1 Introduction

New advances in biology and breakthroughs in genetic research have presented the biotechnology and pharmaceutical industry with a host of promising new targets and compounds to treat a range of diseases. However, the drug development process remains underfunded, with investors shifting capital to other sectors due to mediocre returns on perceived high investment risk. A comparison of five-year periods before and after the recent financial crisis (20042008 vs. 20092013) shows that total funding of drug R&D dropped 21%, from $21.5bn to $16.7bn [15]. Between 2004 and 2012, funding for the National Institutes of Health (NIH) declined by 1.8% per year in real terms [16]. Although funding seems to be improving over the past year in response to a number

of prominent biotech initial public offerings, the capital inflows are highly concentrated among a few large deals, and the number of new start-ups is not increasing [17]. In fact, the lack of funding is particularly severe in early-stage development, prior to Phase II clinical trials. For example, between 2004 and 2011, funding for pre-human/preclinical R&D in the pharmaceutical industry declined by 2.3% per year [16]; 2013 saw only 63 first-time Series A financing rounds in biotechnology, almost 30% lower than the peak of 89 in 2006 and the lowest level in a decade [15]; and the number of active U.S. biotech venture capital firms declined from 201 in 2008 to 138 in 2014 [18].

Fernandez, Stein, and Lo [13] have proposed a "megafund" financing approach, applying portfolio theory and securitization techniques to reduce the risk and enhance the expected returns of a group of investments in drug development projects. Unlike a traditional venture capital fund, the megafund issues both equity and debt ("research-backed obligations" or RBOs), and the portfolio of projects—candidate drugs, licensing agreements, and other intellectual property—serve as collateral for the RBOs. This approach diversifies the typically binary drug investment results across a portfolio of therapeutics, smoothing the portfolio's payout and reducing the volatility of its returns. Securitization also changes the way that cash flows are distributed from a pool of biomedical projects, allowing a broader array of investors to participate in the risk and expected return of drug development according to their risk appetite.

However, issuing securitized debt generally requires collateral that generates a reliable and well-understood stream of cash flows such as an approved drug. Investments in early-stage biomedical projects usually yield no cash flow until they reach Phase IIb, and even then, they provide cash only sporadically, e.g., when they are out-licensed or sold. The unpredictability of both the amount and timing of these cash flows suggests that the megafund is impractical for portfolios exclusively focused on early-stage drug discovery and development.

In this chapter, we extend the concept of the megafund to allow for time-varying amounts of debt or "dynamic leverage," which can accommodate the startup phase of a fund focused purely on preclinical research and development and early-stage translational medicine. Dynamic leverage adjusts the amount of debt that a securitization vehicle can sustain, based on parameters related to its default probability (the likelihood of the entity being unable to meet its payment obligations on a timely basis). It is directly tied to a second concept, "dynamic risk measurement," in which the default risk of a bond is periodically measured via certain credit metrics and performance indicators. Together, dynamic risk measurement and dynamic leverage allow us to con-

struct a time-varying securitization structure that reflects the evolving nature of the portfolio's assets and optimizes the fund's capital structure accordingly.

■ 2.2 Dynamic Leverage

Dynamic leverage is motivated by a simple observation: as a portfolio of biomedical projects progresses, its risk should decrease. Therefore, the amount of debt of a given default probability that can be supported by this portfolio, as a percentage of the total invested capital required, should increase, effectively decreasing the amount of equity required. Because cost of debt (assumed to be 5%–8% here) is lower than cost of equity (usually in the 15% to 30% range), the substitution of equity by debt yields an increase in return on equity. This default probability corresponds to a rating by a Nationally Recognized Statistical Organization (NRSO) such as Moodys Investors Service or Standard & Poor's. The default probability is also referred to as a solvency standard, while the debt as a percent of capital is referred to as an attachment point.

At any point during the life of the fund, there is solvency risk, the risk that the vehicle has insufficient cash to make scheduled interest and/or principal payments. For each rating category, there is an associated solvency standard that specifies the maximum acceptable risk of insolvency (default probability) for that rating class until the notes are repaid (see Table 2.1, provided by Moody's Investors Service [19]). The risk is calculated by examining all of the potential outcomes, and determining what percentage of these outcomes result in an insolvency event. Therefore, the risk is related to a measure of the volatility of future cash flows. The solvency standards are tabulated in Table 2.1 for investment-grade rating categories and several maturities, as published by Moody's Investors Service.

For any given rating tranche, the volatility of the corresponding cash flows may change over time, and therefore the insolvency risk may change. Two factors determine the potential for change in insolvency risk. The primary factor is whether the drug development process is proceeding in accordance to an expected plan (or to the mean of all possible outcomes) at each time instant. If the performance is ahead of the plan, then the probability of insolvency should be lower than the assumed value. In fact, if the performance is on the plan, then the probability should be lower as well because the dispersion of future paths has narrowed, lowering the effective volatility. The second factor is the possibility that inherent volatility has increased due to changes in external factors, the environment, or improved data and forecasts. However, this class

Table 2.1. Moody's "idealized" default probabilities (listed in percentage points) for investment-grade ratings and several maturities.

	Year					
Rating	1	2	3	4	5	6
Aaa	0.0001	0.0002	0.0007	0.0018	0.0029	0.0040
Aa1	0.0006	0.0030	0.0100	0.0210	0.0310	0.0420
Aa2	0.0014	0.0080	0.0260	0.0470	0.0680	0.0890
Aa3	0.0030	0.0190	0.0590	0.1010	0.1420	0.1830
A1	0.0058	0.0370	0.1170	0.1890	0.2610	0.3300
A2	0.0109	0.0700	0.2220	0.3450	0.4670	0.5830
A3	0.0389	0.1500	0.3600	0.5400	0.7300	0.9100
Baa1	0.0900	0.2800	0.5600	0.8300	1.1000	1.3700
Baa2	0.1700	0.4700	0.8300	1.2000	1.5800	1.9700
Baa3	0.4200	1.0500	1.7100	2.3800	3.0500	3.7000

of exogenous events is outside the scope of this dissertation.

Dynamic measurement can be made more precise by employing adaptive trials, during which the posterior probability of success is continuously updated; hence, the amount of debt can be adjusted accordingly. However, for simplicity, we do not use adaptive clinical trials in our model. Dynamic risk measurement is not only useful in determining dynamic leverage, but in any application in which changes in risk have a material impact. For example, in a financing structure that employs guarantees, the guarantee fee can be adjusted dynamically based on the risk profile of the portfolio over time.

In the following section, we present an illustrative example to shed more light on the dynamic leverage concept.

■ 2.2.1 An Illustrative Example for Dynamic Leverage

If we denote the cash inflow of the fund during each period i by C_i, and the scheduled bond principal and interest payments for the end of the same period by P_i, considering that any missed payments constitutes a default event, a default occurs if and only if

$C_i < P_i$ for some i between 1 and n, where n denotes the number of periods in the life of the fund. For example, if each period is six months long and the fund is active for 6 years, we have $n = 12$.

The solvency standard for the rating category of the bond specifies a maximum probability of default for the bond, represented by p_d, which is defined by the following equation:

$$\Pr_0 (C_i \geq P_i,\ 1 \leq i \leq n) \geq 1 - p_d, \tag{2.1}$$

where $\Pr_0(\cdot)$ denotes the probability of an event and the subscript 0 indicates that the probability is evaluated using the information at the inception of the fund, i.e., year 0. Using available information at fund inception, (2.1) states that the probability of having sufficient cash to fulfill all future debt obligations, namely, to stay solvent, should be at least $1 - p_d$ (ideally, we want to get it equal to this level).

For simplicity, let us suppose that only one of the two following scenarios can occur in period 1. In the first scenario, denoted by $s_1 = 1$, there is not enough cash to make the debt payments scheduled for period 1, namely, $C_1 < P_1$, and a default event occurs. In the second scenario, represented by $s_1 = 2$, the generated cash in period 1 is enough to cover the debt payments, i.e., $C_1 \geq P_1$. If we denote the probabilities of the first and second scenarios in period 1 by $\Pr_0(s_1 = 1)$ and $\Pr_0(s_1 = 2)$, respectively, the probability in (2.1) can then be decomposed into two parts:

$$\Pr_0(C_i \geq P_i,\ 1 \leq i \leq n) = \Pr_0(s_1 = 1) \times 0 + \Pr_0(s_1 = 2)\Pr_1(C_i \geq P_i,\ 2 \leq i \leq n) \tag{2.2}$$

$$= \Pr_0(s_1 = 2)\Pr_1(C_i \geq P_i,\ 2 \leq i \leq n) \geq 1 - p_d \tag{2.3}$$

where, by rearranging the terms, we have:

$$\Pr_1(C_i \geq P_i,\ 2 \leq i \leq n) \geq \frac{(1 - p_d)}{\Pr_0(s_1 = 2)} > 1 - p_d. \tag{2.4}$$

Because $0 < \Pr_0(s_1 = 2) < 1$, (2.4) implies that in period 1, if the second scenario, $s_1 = 2$, holds true, the probability of staying solvent over the remainder of the life of the fund, represented by $\Pr_1(C_i \geq P_i,\ 2 \leq i \leq n)$, is larger than the originally intended probability, $1 - p_d$. For example, if $\Pr_0(s_1 = 2) = 99.5\%$, and $p_d = 1\%$, in the second scenario in period 1, the probability of the fund to stay solvent over its life is 99.5%. Therefore, if no extra debt is borrowed in this scenario in period 1, the probability of default on bonds will be guaranteed to be less than 0.5%, which is half of the initially

Table 2.2. Simulation parameters for orphan drug discovery and development.

Phase	Cost (US$ mm)	Success Rate (%)	Duration (years)	Valuation (US$ mm)
Preclinical	5	69	1.00	7.1
Phase I	5	84	1.66	27.6
Phase II	8	53	2.09	75.6
Phase III	43	74	2.15	321.5
NDA	—	96	0.80	701.9
APP	—	—	—	817.6

Abbreviations. mm: million, NDA: New Drug Approval, APP: Approved.

targeted probability of default allowed by the rating of the bonds, i.e., 1%.

This is not surprising because, in the beginning, due to the existence of many different paths with different statistical characteristics, the probability of default must be higher than in the case where some of the "unfavorable" paths have been eliminated. Therefore, to take advantage of this progress in the cash flow of the fund, we can increase the leverage in the second scenario of period 1 (hence, increase the future debt payments, P_is) to increase the expected probability of default back to its allowed level. By dynamically adjusting the level of debt, we can minimize the cost of financing for the fund in each period.

■ 2.3 Dynamic Leverage for an Orphan Drug Fund

For concreteness, we use the statistical model described in [14] to illustrate dynamic risk measurement and dynamic leverage. The focus of Fagnan *et al.* [14] on orphan drugs targeting rare diseases is particularly well-suited for dynamic leverage because these therapies are relatively new and not likely to be able to generate much cash flow at fund inception. To highlight the role of dynamic leverage, we employ the identical orphan drug parameters as in [14].

Following [13] and [14], a discrete-time finite-state Markov chain is employed to model the evolution of each compound through the development cycle. The assumptions regarding the average cost, success rate, duration, and valuation of each phase are listed in Table 2.2.

Under these assumptions, consider an RBO structure to finance a portfolio of investigational therapeutics through their development cycle. In exchange for a pledge

of the future royalty cash flows, equity and debt investors purchase notes and receive a portion of these cash flow streams.

■ 2.3.1 Cash Flow Waterfall

The key governing components of a securitization indenture, or controlling document, are the rules laid out to regulate the flow of cash into the structure, known as a cash flow waterfall. The waterfall sets the guidelines for how and when cash is allocated, and what events trigger default, asset sales, or cash diversion. The flowchart in Figure 2.1 demonstrates the cash flow waterfall for the drug development megafund. Raising debt and making distributions to the equity investors are performed such that the indenture still passes the overcollateralization test after these actions.

■ 2.4 Performance Results

Our simulated RBO portfolio comprises 9 compounds in the preclinical stage and 10 compounds in the clinical Phase I stage. The employed capital structure is composed of one equity tranche, and two debt tranches, namely, mezzanine and senior tranches. The initial amounts of capital for the equity, mezzanine, and senior tranches are \$373.75mm, \$30mm, and \$25mm, respectively, and the annual coupon rates for the mezzanine and senior debt tranches are 8% and 5%, respectively. The maturity dates for the senior and the mezzanine tranche are 4 and 6 years, respectively, and the outstanding balance of each tranche is paid in four equally sized installments over the two years (four semesters) preceding the maturity dates. After 13 semesters (6.5 years), the portfolio of the remaining compounds is liquidated. Assuming that the drug sale takes a year to settle, the cash proceeds from the sale go to the equity investors in the fifteenth semester. Furthermore, any compound, upon reaching a pre-specified target phase (Phase III in the simulations), gets sold regardless of how far into the life of the fund is.

As the portfolio of compounds progresses and its risk decreases over time, the size of the debt tranches—and therefore, the leverage—can be adjusted to maintain a desired probability of default for each tranche. For simplicity, the tranche size adjustment in the simulations is performed only for the mezzanine tranche, and up until the junior bonds start principal repayment, i.e., until the fourth year. Figure 2.2 illustrates the expected size of each tranche as well as the total capital deployed in the portfolio, from both the equity and bond investors, over time.

Several trends in Figure 2.2 are worth noting. As seen in Table 2.2, the compounds

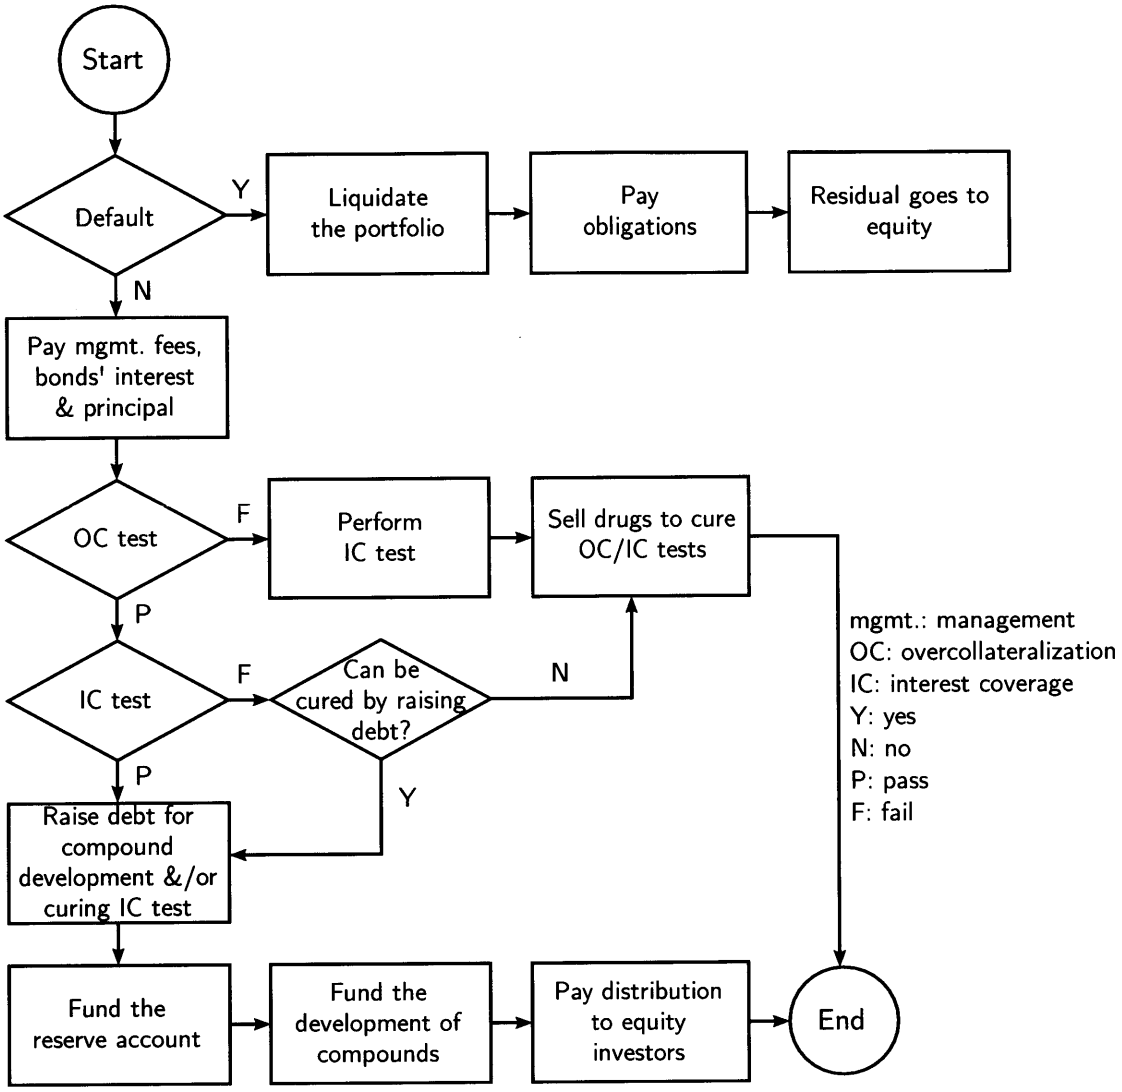

Figure 2.1. Flowchart representation of the cash flow waterfall used in the fund for each six-month period.

need progressively larger amounts of funding as they proceed in their development cycle. If the total required capital is raised in its entirety at the beginning of the fund's life, in anticipation that the compounds will follow their expected path of development, it will impose a drag on the fund's returns. Should this capital be raised by calling more equity, the return on the equity tranche would be diluted. Alternatively, if the financing structure keeps the level of the invested equity constant, issuing more debt at the beginning to meet the expected needs of future drug development, the probability

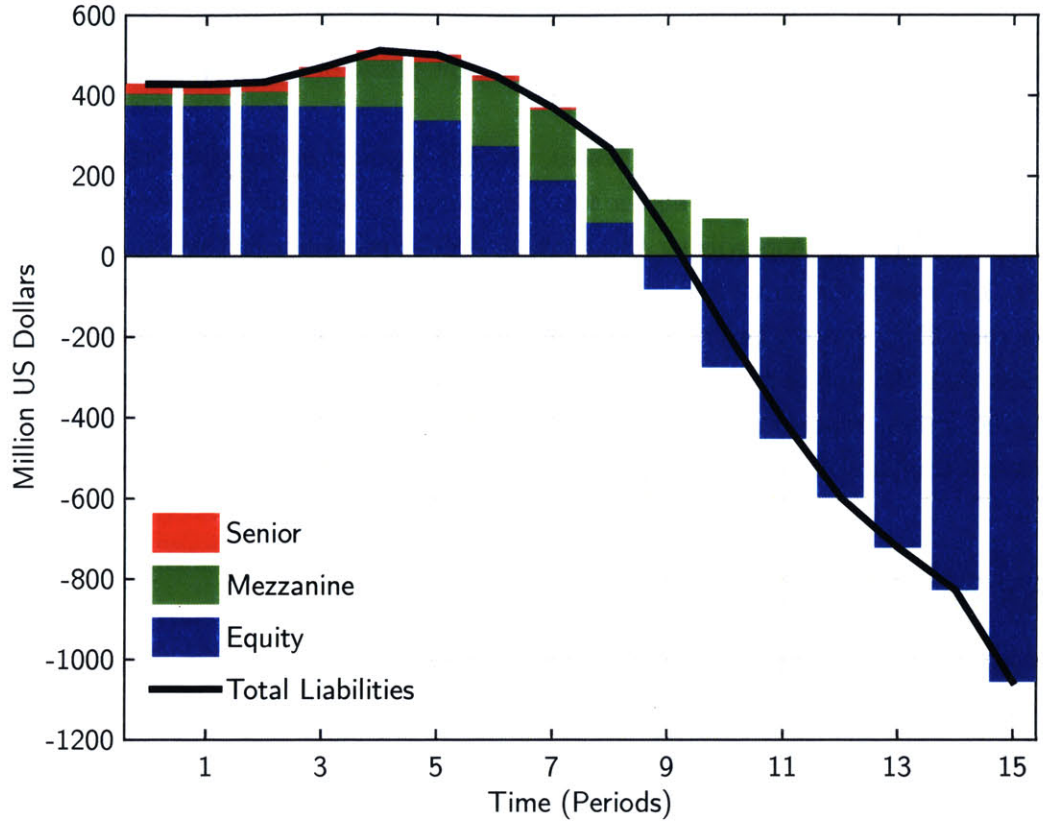

Figure 2.2. Capital structure and total deployed capital in the fund for each six-month period.

of default for the debt tranches would inevitably increase. In this approach, used in [14] , the probability of default increases because the deterioration in portfolio value leads to more debt, while the equity is the same as before. Therefore, the probability of default and the magnitude of loss will increase if more debt is issued at fund inception.

Dynamic leverage can mitigate this issue. More specifically, the mezzanine tranche should increase in size over time to provide the capital required to fund the development of the compounds moving forward in their development cycle. This is done only if raising more debt does not hurt the probability of default for the junior notes; i.e., if it does not increase the solvency risk. Hence, the increase in the mezzanine tranche is slow in earlier periods, when the risk of the portfolio is relatively high, and the debt utilization accelerates as the portfolio moves forward in time and risk is reduced.

A second trend seen in Figure 2.2 relates to the size of the equity tranche, which decreases over time. This is due to distributions made to the equity investors when

the portfolio is on or above the expected path. These distributions come from the sale of those compounds that have reached their target phase of development, and from a portion of the debt raised.

As the risk of the portfolio decreases, we can replace an ever-increasing amount of equity with debt to yield a higher rate of return to the equity investors. This can be achieved without jeopardizing the solvency of the portfolio, as can be observed in Table 2.3, where the simulated expected annualized internal rate of return (IRR) is more than 25%, and the probabilities of default for the senior and mezzanine tranches are less than 0.1 bps and 36.2 bps, respectively. These probabilities of default and the expected losses, reported in Table 2.3, over the life horizon of the senior and junior notes are comparable to that of AAA/Aaa and A+/A1 rated notes, respectively (see Table 2.1).

■ 2.5 Comparison to All-Equity Financing

The third column of Table 2.3, titled "All-EQ 1," compares the RBO structure to an equity structure in which a portfolio of 7 compounds in the preclinical stage and 6 compounds in Phase I is funded using the same level of equity as used in the RBO structure, i.e., $373.75mm. As is observed in Table 2.3, fewer compounds can be financed during the life of the fund under the equity structure compared to the RBO portfolio, since there is no additional injection of capital into the equity portfolio after the initial equity draw. The scientific impact of the equity structure is consequently smaller than that of the RBO portfolio, as measured by the number of the compounds that are sold in Phases II and III. Not only is the scientific impact smaller in the equity structure, but the return characteristics of the equity tranche are not as promising as those of the RBO structure. Due to the debt issuance over time, in the RBO case, more equity is returned to the investors earlier. On the other hand, in the equity structure, the return of capital to the equity investors is constrained by the speed with which the compounds reach the target phase and get sold.

The fourth column in Table 2.3, titled "All-EQ 2," compares the performance of the RBO fund to the performance of the same portfolio of compounds financed by equity alone. The amount of equity used to finance this portfolio is matched to the peak value of the total capital deployed in the RBO structure, i.e., $510.70mm as observed in Figure 2.2. This level is almost 37% more than the RBO's initial equity level of $373.75mm. As is seen in Table 2.3, the scientific impact of this new equity structure

Table 2.3. Comparison of performance results for the RBO portfolio, two equity-financed portfolios, and the static RBO portfolio.

	RBO[a]	All-EQ 1[a,b]	All-EQ 2[a,b]	Static RBO[c]
Number of compounds acquired				
Preclinical	10	7	10	8
Phase I	9	6	9	8
Research impact				
Compounds sold in Phase II	2.6	1.8	2.6	2.2
Compounds sold in Phase III	5.5	3.8	5.5	4.7
Liabilities (US$ millions)				
Capital	428.75	373.75	510.70	575.00
Senior tranche	25.00	—	—	86.25
Initial mezzanine tranche	30.00	—	—	115.00
Equity tranche	373.75	373.75	510.70	373.75
Equity tranche performance				
Expected annualized IRR (%)	25.1	20.7	22.0	13.4
Pr(IRR = −100%) (bps)	38.4	< 0.1	< 0.1	60.0
Pr(IRR < 0%) (%)	10.6	14.5	10.3	13.1
Pr(IRR > 10%) (%)	77.3	69.8	74.6	66.7
Pr(IRR > 25%) (%)	49.8	39.9	42.0	18.4
Debt tranches performance				
Senior tranche				
Probability of default (bps)	< 0.1	—	—	0.8
Expected loss (bps)	< 0.1	—	—	< 0.1
Mezzanine tranche				
Probability of default (bps)	36.2	—	—	56.0
Expected loss (bps)	9.1	—	—	15.0

Abbreviations. RBO: research-backed obligations, IRR: internal rate of return, bps: basis points (1 bp = 0.01%).

[a] All listed numbers are obtained using 20 million Monte Carlo simulation paths for each portfolio.

[b] All-EQ 1 is an equity-financed portfolio where the initial investment is set equal to the initial amount of equity in the RBO portfolio, i.e., $373.75mm, whereas All-EQ 2 is an equity-financed portfolio with the initial investment set to the maximum amount of capital in the RBO portfolio, i.e., $510.70mm (see Figure 2.2).

[c] For static RBO, see Fagnan *et al.* [14].

is the same as that of the RBO structure. However, the financial performance of the equity structure is still less promising than the performance of the RBO since more equity is deployed in the equity structure than in the RBO structure. The only area in which the equity portfolio outperforms the RBO portfolio is the probability of negative returns on the equity. In the equity structure, there is a 10.3% chance of delivering a negative return to the equity investors, whereas this chance is 10.6% for the RBO portfolio, since the equity tranche in the RBO structure is the first to absorb any capital losses. Due to the same reason, the probability that the equity is completely wiped out, i.e., $\Pr(\text{IRR} = -100\%)$, is larger for the dynamic RBO fund (38.4 bps) compared to the all-equity-financed portfolios (< 0.1 bps). However, the upside of the RBO portfolio is much higher than that of the equity portfolios, as measured by the right tail probabilities of their returns reported in Table 2.3, i.e., $\Pr(\text{IRR} > 10\%)$ and $\Pr(\text{IRR} > 25\%)$.

It is clear that adding dynamically leveraged debt to the picture, when feasible and as needed to fund drug development, can enhance both the scientific and the financial impact of the portfolio with little downside risk. Furthermore, if the effect of dynamic leverage were replicated using an equity-financed portfolio, the amount of required equity upfront would be significantly larger (almost 37% more initial equity than the RBO's initial equity as observed in Table 2.3).

■ 2.6 Comparison with Static Capital Structure

For comparison, the performance statistics of the RBO structure with a static capital structure, which was used in [14], are reported in the last column of Table 2.3, labeled "Static RBO". The dynamic RBO clearly outperforms the RBO with a static capital structure from both scientific and financial perspectives. This performance superiority is achieved without jeopardizing the debt performance.

Not only does dynamic leverage increase the return on equity, but it also helps reduce the probability of default for the bondholders in comparison to a static capital structure. This is achieved because less debt is borrowed initially, and more debt issuance happens over time only if the risk of the portfolio permits taking such action. Furthermore, because the probability of default is smaller for this dynamic capital structure than the static RBO used in [14], the volatility of the return on equity is consequently smaller too.

■ 2.7 Robustness Analysis

We check the robustness of our results by varying two key parameters: the value of the approved drug (the bottom right entry in Table 2.2) and the correlation (ρ) of the asset values. Simulations are performed for relative changes in the value of the approved drug from 25% to 10% (denoted by V_{app}) while the correlations of the asset values are varied from 10% to 40%. Different performance measures are then computed—using up to one billion Monte Carlo simulation paths—over the two-dimensional plane constituted by the pairs of V_{app} and ρ drawn from these intervals. The performance measures given in Table 2.3 in the main text thus correspond to a single point on this two-dimensional plane, representing $V_{app} = 0$ and $\rho = 20\%$.

To highlight what may go wrong if the assumptions in the model are too optimistic, the intervals above are not symmetric around the presumed values for the correlation and the approval value. Finally, in our model we use risk-adjusted net present value (rNPV) calculations along with the expected value of an approved orphan drug to derive the expected value of the compound in earlier phases (e.g., Phase I or Phase II). For example, if $V_{app} = -10\%$, not only does this imply a difference between the presumed expected approval value and a typical realized approval value but it also indicates that the values of the compounds in the preclinical stage, Phase I, Phase II, and Phase III are lower than the assumed values in the model. Therefore, the robustness analysis, presented here, simulates the performance of the fund for scenarios where the management team uses these "incorrect" estimates—compared to the market values of the assets—when determining the amount of leverage in each period. This models the situation in which there is a sudden and unexpected event that impacts the market values of the assets in the portfolio.

In the following analysis, four diverse aspects of the fund performance are considered: the statistical characteristics of the return on equity or internal rate of return (IRR), the probability of default (PD) and expected loss (EL) for both debt tranches, and last but not least, the scientific impact of the fund.

When the assets in the fund are mispriced, there will be an inevitable discrepancy between the realized IRR and the presumed IRR (listed in Table 2.3) due to the over/undervaluation of the assets by the fund. However, this mispricing is independent of the amount of the debt involved in the financing of the portfolio's projects. Hence, by changing the approval value, a change in the IRR is expected regardless of how the development of the compounds is financed. On the other hand, it is clear that, if the

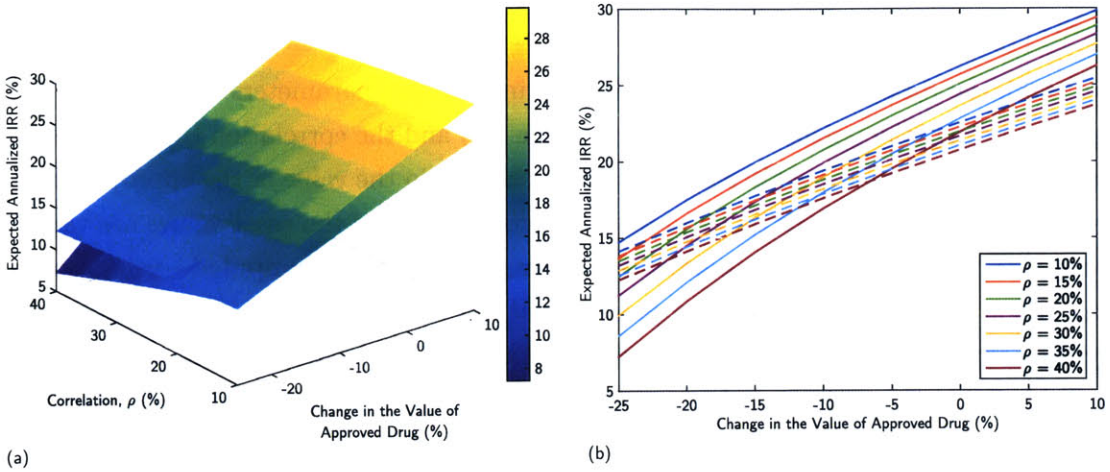

Figure 2.3. Expected annualized IRR for the RBO and equity structures.

Note. The textured shading in (a) and the solid lines in (b) correspond to the RBO structure.

fund's assets are assigned values higher than their market values, more debt will be consequently borrowed than what can be afforded in the capital structure calculations, which in turn increases the PD, and changes the statistical characteristics of the IRR.

To distinguish the role of under-/over-borrowing from the role that the mispricing plays, the RBO performance is, whenever applicable, compared to that of the all-equity-financed portfolio with the initial equity draw of \$510.70mm (All-EQ 2 in Table 2.3). In all of the three-dimensional figures in this section, the textured surfaces correspond to the RBO portfolio, whereas the smooth shading is associated with the all-equity portfolio. In the two-dimensional figures that follow, the solid and dashed lines correspond to the RBO and equity structures, respectively, and each correlation value is represented using a unique color.

Figure 2.3a illustrates the IRR for the RBO and equity portfolios, and Figure 2.3b shows the cross-sections of the surfaces in Figure 2.3a along the V_{app} axis, each associated with a different correlation value. Not surprisingly, both structures deliver higher equity returns for higher asset values. However, there are a few differences worth noting. First, the sensitivity of the IRR in the RBO structure with respect to the approval value (as measured by the line slopes in Figure 2.3b) and the correlation of asset values (as measured by the gap observed between two adjacent lines) is higher than that of the IRR in the equity structure. This higher sensitivity is due to the debt element of the RBO structure. In particular, as the correlation increases and/or the asset values

become smaller, the RBO portfolio becomes less solvent and the sensitivity of its IRR to both ρ and the approval value in turn increases. Therefore, the gap between every two consecutive solid lines is wider on the left side of Figure 2.3b, decreasing as it moves to the right. Furthermore, the slopes of the solid lines in Figure 2.3b are larger on the left side of the figure than the slopes of the same lines on the right side. On the other hand, the slope of the dashed lines and the gaps between them do not quite change, implying a relatively constant sensitivity to the correlation and approval value for the IRR of the equity structure.

Second, as seen in Figure 2.3a, the expected IRR performance of the RBO structure is better than that of the equity structure over most of the $\rho - V_{app}$ plane, except the region for which the correlation and the asset valuations work against the solvency of the portfolio. Specifically, for large correlations, the designed portfolio is not diversified enough, while low asset values make the collateral less valuable. For the worst pair at $\rho = 40\%, V_{app} = -25\%$, the expected IRR for the RBO structure is 7.2%.

It should be noted that the expected IRR for the equity structure would not depend on the correlation among the asset values if there were no distributions paid to the equity investors before the final semester; i.e., all the dashed lines in Figure 2.3b would collapse onto a single line. However, there is a slight dependence of the IRR on the correlation in the equity-financed case in Figure 2.3b, since the equity capital is returned to the investors upon receiving cash from the sale of the compounds reaching the target phase.

The four panels in Figure 2.4 depict the probabilities of large returns (top panels) and negative returns (bottom panels) for the RBO and equity portfolios. Figures 2.4a and 2.4b demonstrate that the IRR distribution of the RBO structure clearly benefits from a larger positive tail relative to the equity structure for all the asset values and correlations considered here. As seen in Figure 2.4b, if the approval value is low, namely less than 80% of the presumed value, the probability of receiving large equity returns for the equity structure is the highest for the largest correlation value. This is in sharp contrast with the trend for higher approval values, since the probability of large equity returns is inversely proportional to the correlation for higher approval values. This is an immediate impact of the correlation on the tail of the IRR distribution. Intuitively, if the asset values are small, then a relatively large correlation is required to yield a large return on equity.

While the RBO structure enjoys an IRR distribution with a large positive tail, the negative tail of the IRR distribution is almost the same for both structures. In

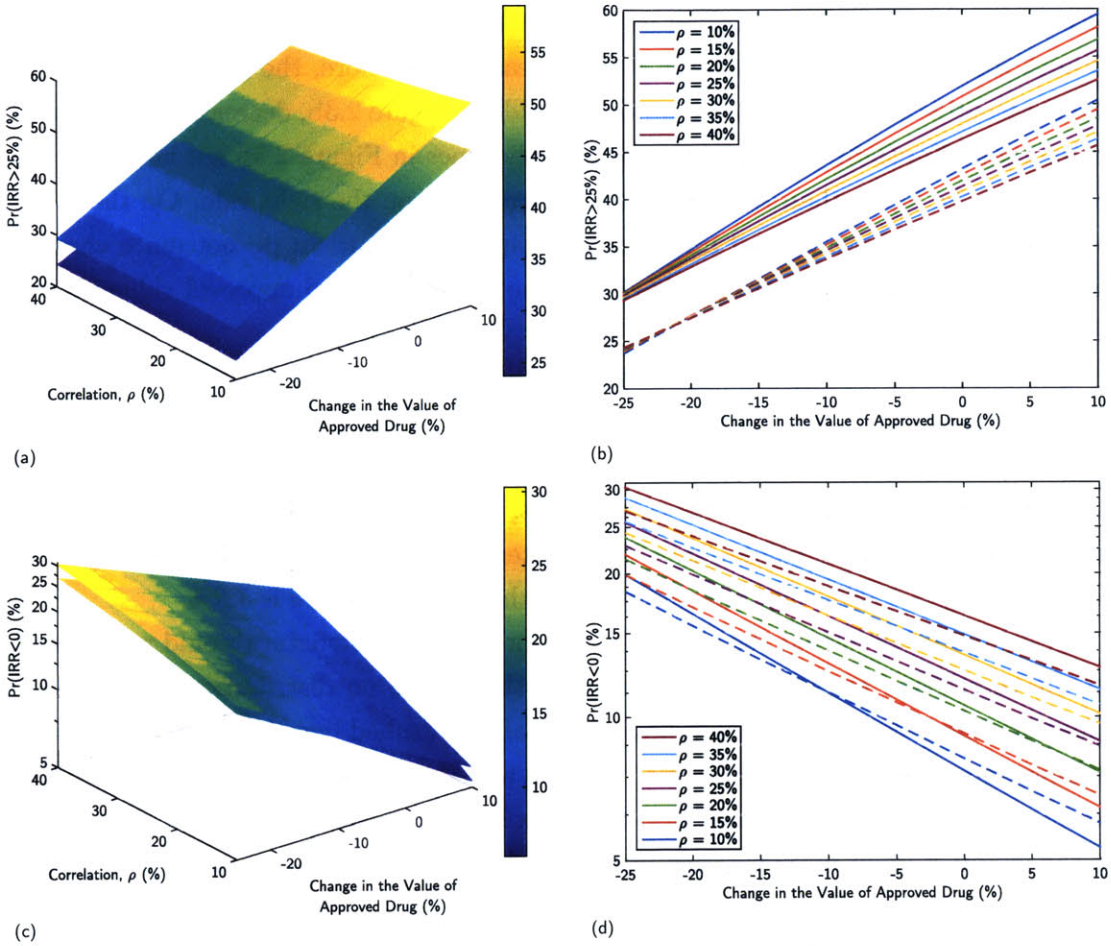

Figure 2.4. The probability of yielding large returns for the RBO and equity structures is illustrated in panels (a) and (b) while the probability of delivering negative returns on these portfolios is presented in (c) and (d).

Note. The textured shading in (a) and (c) and the solid lines in (b) and (d) correspond to the RBO structure.

particular, regarding the probability of negative equity returns, for each correlation value, there exists an approval value below which the RBO performance is worse than the equity performance, and above which the RBO probability of a negative IRR is less than that of the equity structure. This inflection point is seen in 2.4d for correlations up to 20%. If the correlation value is larger than 20%, then for all the approval values considered here, the RBO structure has a larger probability of yielding negative returns than the equity structure, due to the fact that for large correlation values, the RBO

portfolio is not diversified enough, becoming vulnerable to over-borrowed debt.

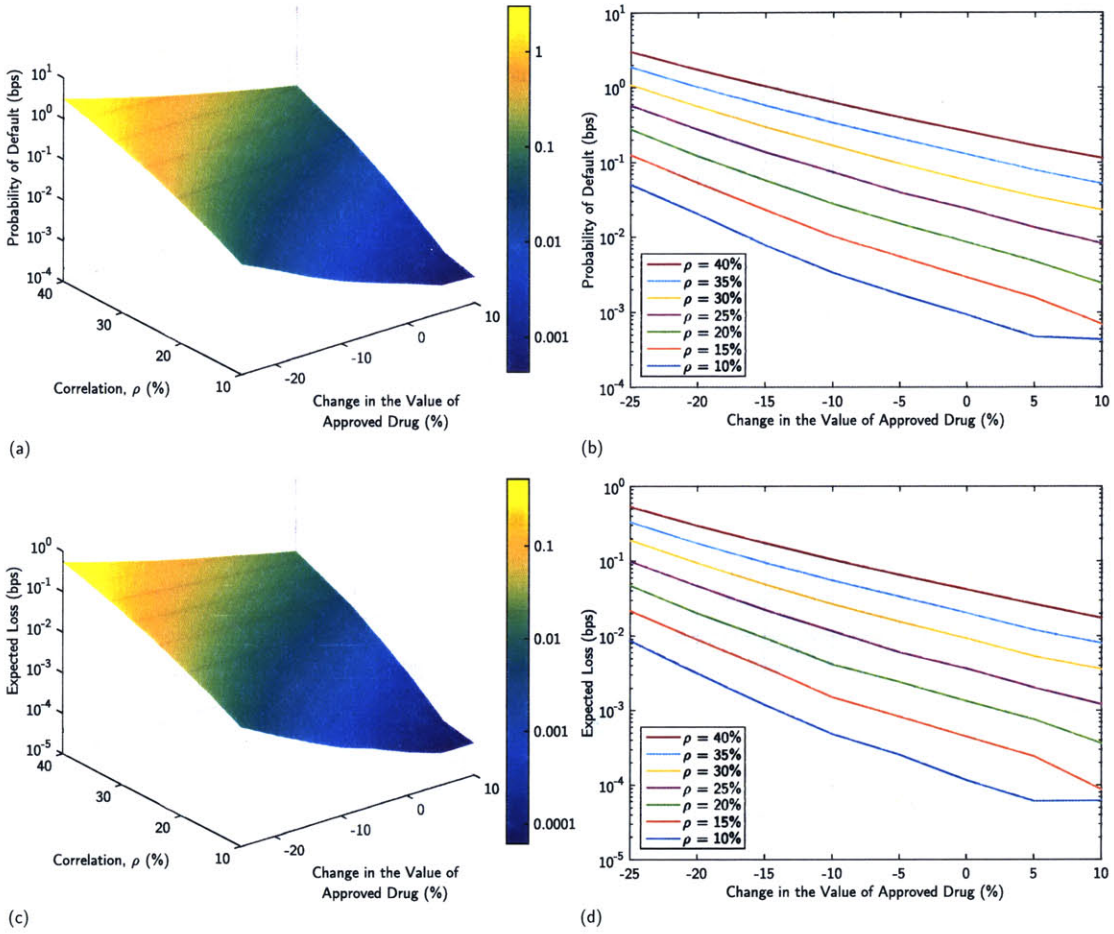

Figure 2.5. Probability of default for the senior tranche is illustrated in panels (a) and (b) and the expected loss for the senior tranche is presented in (c) and (d).

Figures 2.5 and 2.6 depict the PDs and ELs for the senior notes and the junior notes, respectively. For the senior and mezzanine tranches, the PD and EL increase with the correlation of asset values and with a decrease in the approval value. The worst-case scenario for the senior tranche is for $\rho = 40\%$ and $V_{app} = -25\%$, where PD = 3.0 bps and EL = 0.5 bps, respectively. For the junior notes, the worst-case scenario corresponds to the same scenario as in the senior tranche, for which PD = 6.6% and EL = 2.6%.

As with the previous performance criteria, it is clear in Figure 2.7 that the number of compounds sold in Phases II and III by the RBO structure is more sensitive to

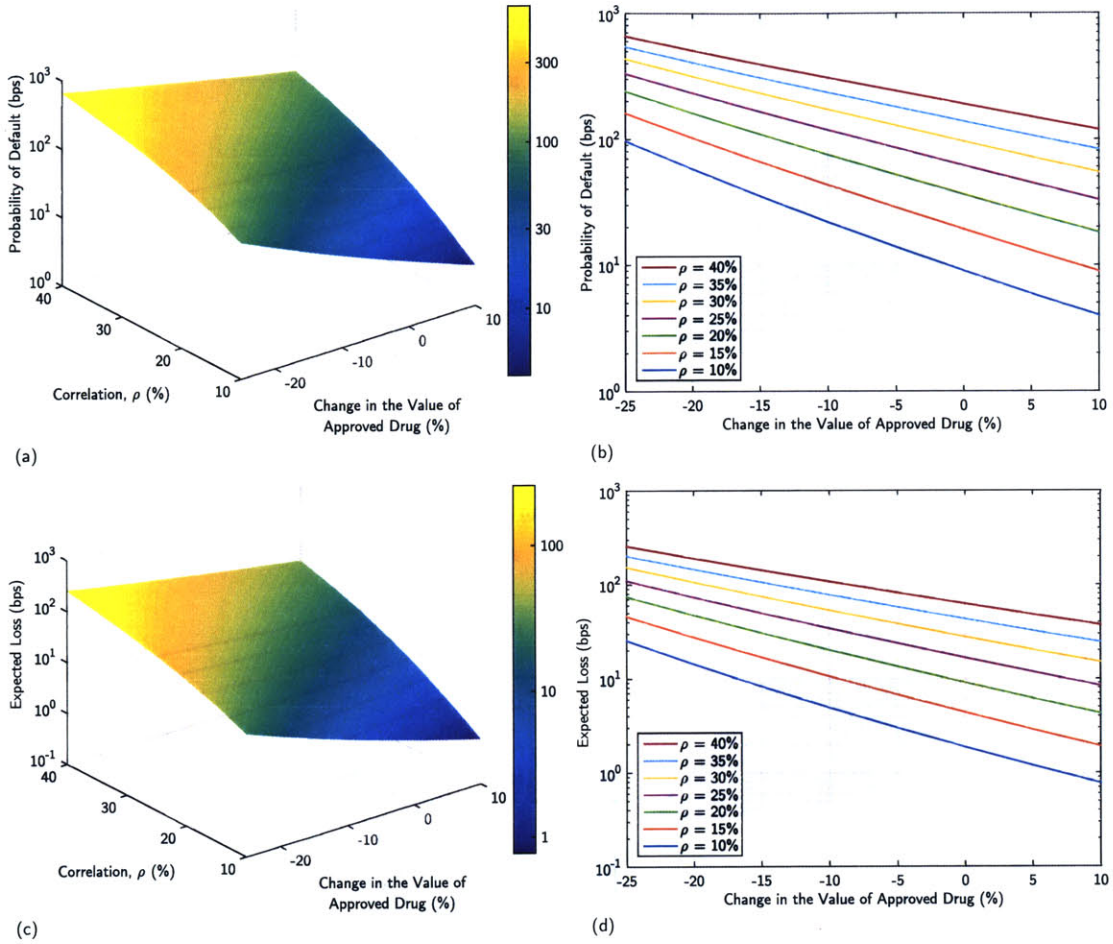

Figure 2.6. Probability of default for the mezzanine tranche is illustrated in panels (a) and (b) and the expected loss for the mezzanine tranche is presented in (c) and (d).

both the change in the approval value and the correlation of asset values than is the corresponding number of compounds sold by the equity structure. Ideally, for the best scientific impact, the average performance of the equity structure should be completely independent of the correlation. However, there is a slight dependence on ρ, as seen in Figures 2.7b and 2.7d, because of distributions made to the equity investors.

■ 2.8 Conclusion

The application of portfolio theory and securitization techniques to financing drug development has the potential to be a disruptive technology. In this chapter we propose a more efficient structure and higher returns to equity for investors by adding dy-

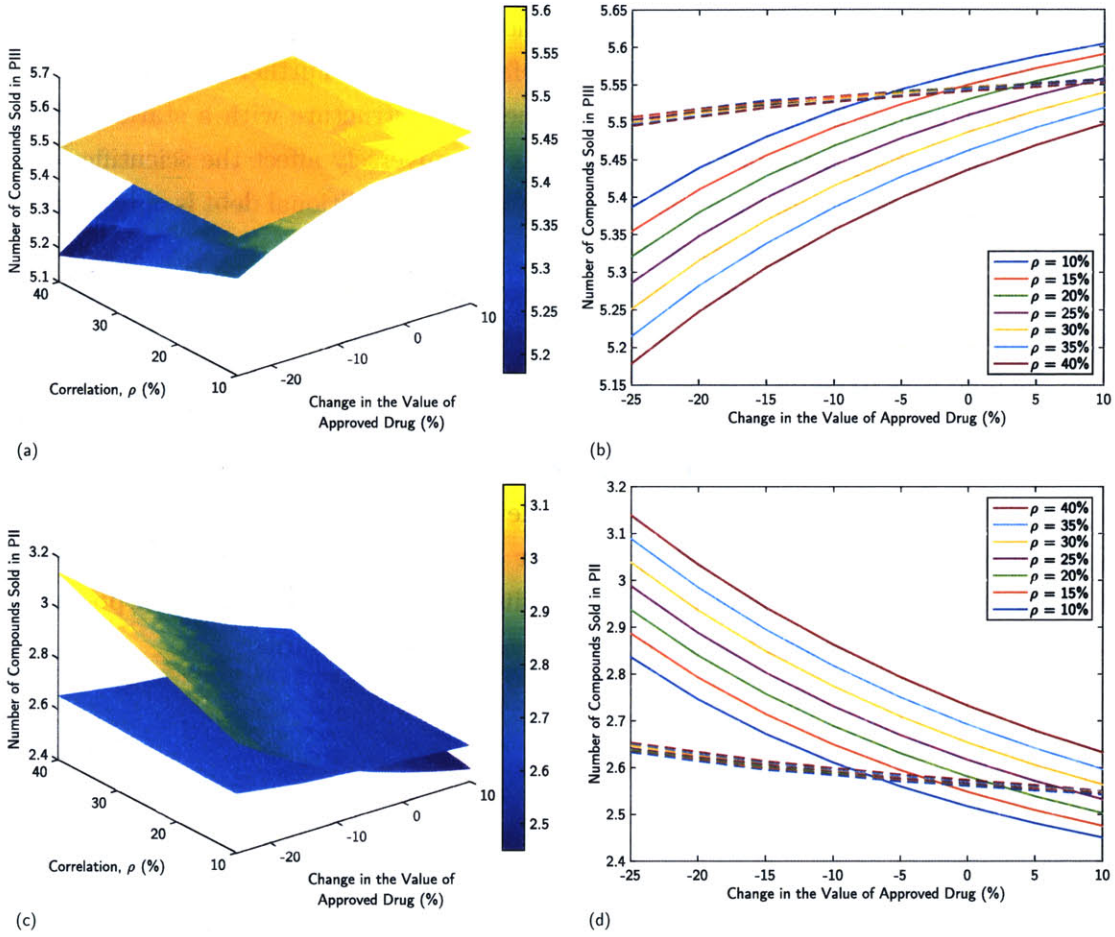

Figure 2.7. The number of compounds sold in Phase III by the RBO and equity structures is illustrated in panels (a) and (b) and the number of compounds sold in Phase II in these portfolios is presented in (c) and (d).

Note. The textured shading in (a) and (c) and the solid lines in (b) and (d) correspond to the RBO structure.

namic leverage, a novel securitization technique, to the megafund structure proposed in [13, 14]. There are, of course, a number of practical challenges to launching and managing a megafund. A comprehensive discussion of these challenges is beyond the scope of this dissertation, but we address some of the most pressing issues in Appendix A such as how the fund would be managed, whether the parameters we have assumed are realistic, and how. Several other recent studies offer more detailed analysis of these challenges and how they can be addressed [20–25]. The main result of this study is that a fund incorporating dynamic leverage requires less upfront equity to finance the

development of the compounds in the portfolio than previous implementations, and generates higher returns with similar risks of default and loss. Furthermore, the volatility of equity returns is lower compared to a megafund structure with a static capital structure. Borrowing more debt over time does not adversely affect the scientific outcome because in the dynamically leveraged approach, the additional debt is only needed if the portfolio is on its expected path.

Dynamic leverage magnifies performance, both positive and negative. If the actual performance of the portfolio of projects is better than indicated by prior assumptions, then the fund with dynamic leverage will outperform an equity-financed portfolio. If the portfolio underperforms, however, then the equity-funded portfolio will perform better. This result is expected, given the nature of leverage. The higher volatility (risk) of equity returns in a megafund with dynamic leverage, as compared to an all-equity-financed portfolio, is accompanied by a higher expected equity return. Nevertheless, if further securitization technologies are introduced into the pharmaceutical portfolio structure, we expect commensurate improvements to equity returns.

Chapter 3

Drug Approval

Implicit in the drug-approval process is a trade-off between Type I and Type II error. We propose using Bayesian decision analysis (BDA) to minimize the expected cost of drug approval, where relative costs are calibrated using U.S. Burden of Disease Study 2010 data. The results for conventional fixed-sample randomized clinical-trial designs suggest that for terminal illnesses with no existing therapies such as pancreatic cancer, the standard threshold of 2.5% is too conservative; the BDA-optimal threshold is 27.9%. However, for relatively less deadly conditions such as prostate cancer, 2.5% may be too risk-tolerant or aggressive; the BDA-optimal threshold is 1.2%. We compute BDA-optimal sizes for 25 of the most lethal diseases and show how a BDA-informed approval process can incorporate all stakeholders' views in a systematic, transparent, internally consistent, and repeatable manner.

■ 3.1 Introduction

Randomized clinical trials (RCTs) have been widely accepted as the most reliable approach for determining the safety and efficacy of drugs and medical devices [26, 27], and their outcomes largely determine whether new therapeutics are approved by regulatory agencies such as the U.S. Food and Drug Administration (FDA). Because RCTs often involve several thousand human subjects and require years to complete, the FDA is sometimes criticized for being too conservative, requiring trials that are "overly large" [28] and using too conservative a threshold of statistical significance.

In response to these concerns, the FDA has gone to great lengths to expedite the approval process for drugs intended to treat serious conditions and rare diseases [29, 30].[1] Four programs—fast-track, breakthrough-therapy, accelerated-approval, and priority-review designations—provide faster reviews and/or use surrogate endpoints to judge

[1] See `http://www.fda.gov/forpatients/approvals/fast/ucm20041766.htm`

efficacy. However, the published descriptions [29, 30] do not indicate any difference in the statistical thresholds used in these programs versus the standard approval process, nor do they mention adapting these thresholds to the severity of the disease. Hence, from the patient's perspective, the approval criteria in these programs may still seem too conservative, especially for terminal illnesses with no existing treatment options. Moreover, a large number of compounds are not eligible for these special designations, and some physicians have argued that the regulatory safety requirements for drugs targeting non-cancer life-threatening diseases, e.g., cirrhosis of the liver and hypertensive heart disease, should be relaxed.

At the heart of this debate is the unavoidable regulatory trade-off between maximizing the benefits of effective therapies to patients and minimizing the risk to those who do not respond to such therapies. Even under the current thresholds of statistical significance, both the U.S. and Europe have seen harmful drugs with severe side effects make their way into the market [31–34]. Therefore, the FDA and the European Medicine Agency (EMA)—government agencies mandated to protect the public—are understandably reluctant to employ more risk-tolerant or aggressive statistical criteria to judge the efficacy of a drug. However, we show in this chapter that when the risk of adverse side effects is explicitly weighed against the severity of the disease, the standard thresholds of statistical significance are often too conservative for the most serious afflictions such as pancreatic cancer. On the other hand, the same conventional statistical thresholds can be too aggressive for milder illnesses such as prostate cancer. Therefore, criticizing drug regulatory agencies for being overly conservative or aggressive without explicitly specifying the burden of disease, i.e., the therapeutic costs and benefits for current and future patients, is uninformed and vacuous.

In statistical terms, regulators must weigh the cost of a Type I error—approving an ineffective therapy—against the cost of a Type II error—rejecting an effective therapy. However, the term "cost" in this context refers not just to direct financial costs, but also includes the consequences of incorrect decisions for all current and future patients. Complicating this process is the fact that these trade-offs sometimes involve utilitarian conundrums in which small benefits for a large number of patients must be weighed against devastating consequences for an unfortunate few. Moreover, the relative costs (risks) of the potential outcomes are viewed quite differently by different stakeholders; patients dying of pancreatic cancer may not be as concerned about the dangerous side effects of an experimental drug as a publicly traded pharmaceutical company whose shareholders will bear the enormous cost of wrongful death litigation.

The need to balance these competing considerations in decision-making for drug-approval has long been recognized by clinicians, drug-regulatory experts and other stakeholders [35–37]. It has also been recognized that these competing factors should be taken into account when designing clinical trials [38–40] and one approach to quantify this need is to assign different costs to the different outcomes [40].

In this chapter, we propose to make these trade-offs explicit by applying a Bayesian decision analysis (BDA) framework to the design of RCTs as advocated by [40, 41]. In this framework, Type I and II errors are assigned different costs, as first suggested by [38–40], but we also take into account the delicate balance between the costs associated with an ineffective treatment during and after the trial. Given these costs, other population parameters, and prior probabilities, we can compute an expected cost for any fixed-sample clinical trial and minimize the expected cost over all fixed-sample tests to yield the BDA-optimal fixed-sample trial design.

The concept of assigning costs to outcomes and employing cost-minimization techniques to determine optimal decisions is well known [42]. Our main contribution is to apply this standard framework to the drug-approval process by explicitly specifying the costs of Type I and Type II errors using burden-of-disease data. This approach yields a systematic, objective, transparent, and repeatable process for making regulatory decisions that reflects differences in disease-specific parameters. Moreover, given a specific statistical threshold, and assuming that this threshold is optimal from a BDA perspective, we can invert the relationship between cost parameters and their corresponding BDA-optimal tests to impute the costs implicit in a given clinical trial design. This allows us to infer the FDA's implicit weighting of Type I and II errors, which yields an objective measure of whether its approval thresholds are too conservative or aggressive.

Using U.S. Burden of Disease Study 2010 data [43], we show that the current standards of drug-approval are weighted more on avoiding a Type I error (approving ineffective therapies) rather than a Type II error (rejecting effective therapies). For example, the standard Type I error of 2.5% is too conservative for clinical trials of therapies for pancreatic cancer—a disease with a 5-year survival rate of 1% for stage IV patients (American Cancer Society estimate, last updated 3 February 2013) [44]. The BDA-optimal size for these clinical trials is 27.9%, reflecting the fact that, for these desperate patients, the cost of trying an ineffective drug is considerably less than the cost of not trying an effective one. On the other hand, 2.5% may be too aggressive for clinical trials testing prostate cancer therapies, for which the BDA-optimal significance level is 1.2%. It is worth noting that the BDA-optimal size is larger not just for life-threatening

cancers but also for serious non-cancer conditions, e.g., cirrhosis of the liver (optimal size = 16.6%) and hypertensive heart disease (optimal size = 8.1%).

Although there are obvious utilitarian reasons for weighting Type I errors more heavily, they do not necessarily apply to all diseases or stakeholders. For terminal illnesses where patients have no choice but death, the relative costs of Type I and II errors are very different than for non-life-threatening conditions. This difference is clearly echoed in the Citizens Council report published by the U.K.'s National Institute for Health and Care Excellence (NICE) [45], and has also been documented in a series of public meetings held by the FDA as part of its five-year Patient-Focused Drug Development Program, in which the gap between patients' risk/benefit perception and the FDA's was apparent [46, 47]. Our BDA framework incorporates the severity of the disease into its design—as advocated in part 312, subpart E of title 21 Code of Federal Regulation (CFR) [48]—and the FDA reports [46, 47], among many other sources, can be used to determine the relative cost parameters from the patients' and even the general public's perspective in an objective and transparent manner. As suggested in [49], using hard evidence, i.e., available data, for assigning costs to different events is a feasible remedy to the controversy often surrounding Bayesian techniques due to their subjective judgment factor in the cost-assignment process. In fact, Bayesian techniques have survived controversy and are currently used extensively in clinical trials for medical devices, mainly due to the support received from the FDA's Center for Devices and Radiological Health (CDRH) and the use of hard evidence in forming priors in those trials [49].

In Section 3.2, we describe the shortcomings of a classical approach in designing a fixed-sample test. We then lay out the assumptions about the clinical trial to be designed, and the primary response variable affected by the drug in Section 3.3. The BDA framework is introduced in Section 3.4, which can be shown to mitigate the shortcomings of the classical approach, and the BDA-optimal fixed-sample test is then derived. We apply this framework in Section 3.5 by first estimating the parameters of the Bayesian model using the U.S. Burden of Disease Study 2010 [43]. Using these estimates, we compute the BDA-optimal tests for 25 of the top 30 leading causes of death in the U.S. in 2010 and report the results in Section 3.6. We conclude in Section 3.7.

■ 3.2 Limitations of the Classical Approach

Two objectives must be met when determining the sample size and critical value for any fixed-sample RCT: (1) the chance of approving an ineffective treatment should be minimized; and (2) the chance of approving an effective drug should be maximized. The need for maximizing the approval probability for an effective drug is obvious. In the classical (frequentist) approach to hypothesis testing—currently the standard framework for designing clinical trials—these two objectives are pursued by controlling the probabilities of Type I and Type II errors. Type I error occurs when an ineffective drug is approved, and the likelihood of this error is usually referred to as the size of the test. Type II error occurs when an effective drug is rejected, and the complement of the probability of this error is defined as the power of the test.

It is clear that, for a given sample size, minimizing one of these two error probabilities is in conflict with minimizing the other (for example, the probability of a Type I error can be reduced to 0 by rejecting all drugs). Therefore, a balance must be struck between them. The classical approach addresses this issue by constraining the probability of Type I error to be less than a fixed value, usually $\alpha = 2.5\%$ for one-sided tests, and, by choosing a large enough sample size, it maintains a power for the alternative hypothesis, right around another somewhat arbitrary level, usually $1 - \beta = 80\%$.

The arbitrary nature of these values for the size and power of the test raises legitimate questions about their justification. As will be seen later, these particular values correspond to a specific situation, which need not (and most likely does not) apply to clinical trials employed to test new drugs for different diseases. It is also worth noting that these numbers were brought to the design paradigm of clinical trials from other industries, in particular, the manufacturing industry. Therefore, it is reasonable to ask if these totally different industries should use the same values for the size and power of their tests. The consequences of wrongly rejecting a high-quality product in quality testing must be much different from the results of mistakenly rejecting an effective drug for many patients with a life-threatening disease, who may desperately be looking for effective therapeutics. In other words, there must be different *costs* associated with each of these wrong rejections.

In addition to the arbitrary nature of the commonly used values for the size and power of tests, there is an important ethical issue with regard to the classical design of clinical trials. The frequentist approach aims to minimize the chance of ineffective treatment after the trial, which is caused by Type I error. However, it does not take into

account the ineffective treatment *during* the trial, and dismisses that *at least* half of the recruited subjects are exposed to ineffective treatment during the trial, assuming a balanced two-arm RCT [40, 50]. This ethical issue, along with financial considerations, is the principal reason that the sample size in classical trial design is not further increased to get more power. Recently there have been more novel frequentist designs for clinical trials, e.g., group sequential and adaptive tests, to decrease the average sample size in order to mitigate this ethical issue. However, one shortcoming of all these approaches is that they do not take into account the severity of the target disease.

Finally, the classical approach to the design of clinical trials does not take into account the possible number of patients who will eventually be affected by the outcome of the trial. Patients suffering from the target disease may be affected positively in the case of an approved effective drug, or adversely in the case of an approved ineffective drug or a rejected effective drug. From this and similar arguments, it is clear that the sample size of the trial should depend on the size of the population of patients who will be affected by the outcome of the trial, as suggested in [39, 50, 51]. We refer to the population to be affected by the outcome of the trial as *the target population* in the rest of this chapter, and note that it is the same as the patient horizon originally proposed in [38, 39] and later used in [50, 51]. This idea has an immediate and intuitive consequence: If the target population of a new drug comprises 100,000 individuals, its clinical trial must be larger than a trial designed for a drug with a target population of only 10,000 individuals.

■ 3.3 A Review of RCT Statistics

In this section, we explain the basic statistics of RCTs and define the notation employed in this chapter. We begin with the design of the balanced two-arm RCT where the subjects are randomly assigned to either the treatment or control arm, and there is an equal number of subjects in each arm. For simplicity, the focus is only on fixed-sample tests, where the number of subjects per arm, denoted by n, is determined prior to the trial and before making any observations. Furthermore, only after collecting *all* the observations, shall a decision be made on whether or not the drug is effective. However, our approach is equally applicable to more sophisticated designs since the more novel designs usually try to mimic the statistical performance of a fixed-sample test, e.g., frequentist power and size, while minimizing sample size.

A quantitative primary endpoint is assumed for the trial. For instance, the endpoint

may be the level of a particular biochemical in the patient's blood, which is measured on a continuous scale and modeled as a normal random variable [27, 52]. The subjects in the treatment and control arms receive the drug and placebo, respectively, and each subject's response is independent of all other responses. It is worth noting that if there exists a current treatment in the market for the target disease of the drug, then the existing drug, instead of the placebo, is assumed to be administered to the patients in the control arm. In either situation, it is natural to assume that the administered drug to the control arm patients is not toxic. The response variables in the treatment arm, denoted by $\{T_1, \ldots, T_n\}$, are independent and identically distributed (iid), where $T_i \overset{\text{iid}}{\sim} \mathcal{N}(\mu_t, \sigma^2)$. Similarly, for the control (placebo) arm responses, represented by $\{X_1, \ldots, X_n\}$, we assume $X_i \overset{\text{iid}}{\sim} \mathcal{N}(\mu_x, \sigma^2)$, where the response variance in each arm is known and equal to σ^2. The response variance is assumed to be the same for both arms, but this assumption can easily be relaxed.

Furthermore, we focus only on superiority trials, in which the drug candidate is likely to have either a positive effect or no effect (possibly with adverse side effects).[2] Let us define the treatment effect of the drug, δ, as the difference of the response means in the two arms, i.e., $\delta \triangleq \mu_t - \mu_x$. The event in which the drug is ineffective and has adverse side effects defines our null hypothesis, H_0, corresponding to $\delta = 0$ (and the assumption of side effects is meant to represent a "worst-case" scenario since ineffective drugs need not have any side effects). On the other hand, the alternative hypothesis, H_1, represents a positive treatment effect, $\delta = \delta_0 > 0$. Therefore, a one-sided superiority test is appropriate for distinguishing between these two point hypotheses.

In a fixed-sample test with n subjects in each arm, we collect observations from the treatment and control arms, namely, $\{T_i\}_{i=1}^n$ and $\{X_i\}_{i=1}^n$, respectively, and form the following Z-statistic (sometimes referred to as the Wald statistic):

$$Z_n = \frac{\sqrt{\mathcal{I}_n}}{n} \sum_{i=1}^{n} (T_i - X_i), \tag{3.1}$$

where Z_n is a normal random variable, i.e., $Z_n \sim \mathcal{N}(\delta\sqrt{\mathcal{I}_n}, 1)$, and $\mathcal{I}_n = \frac{n}{2\sigma^2}$ is the so-called information in the trial [52]. The Z-statistic, Z_n, is then compared to a critical value, λ_n, and the null hypothesis is not rejected, denoted by $\widehat{H} = H_0$, if the Z-statistic is smaller than the critical value. Otherwise, the null hypothesis is rejected, represented

[2]Non-inferiority trials—where a therapy is tested for similar benefits to the standard of care but with milder side effects—also play an important role in the biopharma industry, and our framework can easily be extended to cover these cases.

by $\widehat{H} = H_1$:

$$Z_n \underset{\widehat{H}=H_0}{\overset{\widehat{H}=H_1}{\gtrless}} \lambda_n. \tag{3.2}$$

As is observed in (3.2), the critical value used to reject the null hypothesis, or equivalently the statistical significance level, is allowed to change with the sample size of the trial, hence the subscript n in λ_n. This lends more flexibility to the trial than the classical setting, where the significance level is exogenous and independent of the sample size. Since a fixed-sample test is completely characterized by two parameters, namely, its sample size and critical value, as seen in (3.2), we denote a fixed-sample test with n subjects in each study arm and a critical value λ_n by $\texttt{fxd}(n, \lambda_n)$. It should be noted that, for the sake of simplicity, we use sample size and number of subjects per arm interchangeably throughout this work. Finally, the assumption that individual response variables are Gaussian is not necessary. Instead, as long as the assumptions of the Central Limit Theorem hold, the distribution of the Z-statistic, Z_n, in (3.1) follows an approximately normal distribution. Therefore, this model should be broadly applicable to a wide range of contexts.

■ 3.4 Bayesian Decision Analysis

The costs associated with a clinical trial can be categorized into two groups: in-trial costs and post-trial costs, where in-trial costs, while independent of the final decision of the clinical trial, depend on the number of subjects recruited in the trial. Post-trial costs, on the other hand, depend solely on the final outcome of the trial and are assumed to be independent of the number of recruited patients. In particular, assume there is no post-trial cost associated with making a correct decision, i.e., rejecting an ineffective drug or approving an effective drug. We further allow asymmetric post-trial costs associated with Type I and Type II errors, denoted by C_1 and C_2, respectively. For brevity, let us call "the post-trial cost associated with Type I error" simply the *Type I cost*, and similarly for the *Type II cost*.

Specifying asymmetric costs for Type I and Type II errors allows us to incorporate the consequences of these two errors with different weights in our formulation. For example, in the case of a life-threatening disease, where patients can benefit tremendously from an effective drug, the Type II cost—caused by mistakenly rejecting an effective drug—must be much larger than the Type I cost, i.e., $C_1 \ll C_2$. On the other hand, if the disease to be treated is mild, e.g., mild anemia or secondary infertility, the cost of

Table 3.1. Post-trial and in-trial costs associated with a balanced fixed-sample randomized clinical trial.

	Post-Trial		In-Trial
	$\widehat{H} = H_0$	$\widehat{H} = H_1$	
$H = H_0$	0	C_1	nc_1
$H = H_1$	C_2	0	$n\gamma C_2$

$C_1 = Nc_1$ and $C_2 = Nc_2$.

adverse side effects can be much larger than the cost of not approving an effective drug for the disease, hence, the Type I cost can be much larger than the Type II cost, i.e., $C_1 \gg C_2$. If the severity of the disease is intermediate, e.g., moderate anemia or mild dementia, then these two post-trial costs may be more or less the same, i.e., $C_1 \approx C_2$.

Furthermore, the two post-trial costs, C_1 and C_2, are assumed to be proportional to the size of the target population of the drug. The larger the prevalence of the disease, the higher the cost caused by a wrong decision in favor of/against the null hypothesis; therefore, the larger the values of C_1 and C_2. Let us assume this relation is linear in the target population size. More precisely, if the size of the target population is N, assume there exist two constants, c_1 and c_2, which are independent of the disease prevalence and depend only on the adverse side effects of the drug and the characteristics of the disease, respectively, such that the following linear relation holds:

$$C_i = Nc_i, \quad i = 1, 2, \tag{3.3}$$

where c_1 and c_2 can be interpreted as the cost per person for Type I and Type II errors, respectively. Lower case letters represent cost per individual, while uppercase letters are used for aggregate costs.

In-trial costs are mainly related to patients' exposure to inferior treatment, e.g., the exposure of enrolled patients to an ineffective but toxic drug in the treatment arm or the delay in treating *all* patients (in the control group and in the general population) with an effective drug. If the drug being tested is ineffective, since there are n subjects in the treatment arm taking this drug, they collectively experience an in-trial cost of nc_1. In this case, the patients in the control arm experience no extra cost, since the current treatment or the placebo is assumed not to be toxic. However, if the drug is effective, the situation is quite different. In this case, for every additional patient in the

trial, there will be an incremental delay in the emergence of the drug in the market. This delay affects all patients, both inside or outside the trial. Therefore, we model this cost to be a fraction of the aggregate Type II cost C_2, and linear in the number of subjects in the trial, n. To be more specific, we assign an in-trial cost of $n\gamma C_2$ for an appropriate choice of γ (for the results presented in Section 3.6, we use $\gamma = 4 \times 10^{-5}$). All the cost categories associated with a fixed-sample test are tabulated in Table 3.1.

For a given fixed-sample test $\texttt{fxd}(n, \lambda_n)$, where Z_n is observed, and the true underlying hypothesis is H, we can define the incurred cost, denoted by $C(H, Z_n, \texttt{fxd}(n, \lambda_n))$, as the following:

$$C(H, Z_n, \texttt{fxd}(n, \lambda_n)) = \begin{cases} Nc_1 \mathbb{1}_{\{Z_n \geq \lambda_n\}} + nc_1, & H = H_0 \\ \\ Nc_2 \mathbb{1}_{\{Z_n < \lambda_n\}} + n\gamma N c_2, & H = H_1 \end{cases}, \tag{3.4}$$

where $\mathbb{1}$ is the indicator function and takes on the value 1 when its argument is true, and is equal to zero otherwise. Here, the first line corresponds to the case where the drug is ineffective, denoted by $H = H_0$. In this case, there will be a post-trial cost, C_1, caused by Type I error, i.e., approving the ineffective drug, which yields the first term. The second term in the first line is the in-trial cost of having n patients in the treatment arm taking this ineffective drug. The second line in (3.4) represents the case, in which the drug is effective, denoted by $H = H_1$. In this case, the second term is the in-trial cost, as explained earlier, and the first term is due to rejecting the effective drug, i.e., if $Z_n < \lambda_n$, resulting in the post-trial Type II cost.

BDA-Optimal Fixed-Sample Test

Let us assume prior probabilities of p_0 and p_1 for the null and alternative hypotheses, respectively, i.e., $P(H_0) = p_0$ and $P(H_1) = p_1$, where $p_0, p_1 > 0$ and $p_0 + p_1 = 1$. It is then straightforward to calculate the expected value of the cost, associated with $\texttt{fxd}(n, \lambda_n)$ and given by (3.4), as the following:

$$\begin{aligned} C(\texttt{fxd}(n, \lambda_n)) &\triangleq E[C(H, Z_n, \texttt{fxd}(n, \lambda_n))] \\ &= p_0 c_1 \left[N\Phi(-\lambda_n) + N\bar{c}_2 \Phi(\lambda_n - \delta_0 \sqrt{\mathcal{I}_n}) + n(1 + \gamma N \bar{c}_2) \right], \end{aligned} \tag{3.5}$$

where Φ is the cumulative distribution function of a standard normal random variable, $Z \sim \mathcal{N}(0, 1)$, and E is the expectation operator. It is worth noting that if $p_0 = p_1 = 0.5$,

then $\overline{c}_2 = \frac{p_1 c_2}{p_0 c_1}$ reduces to $\overline{c}_2 = \frac{c_2}{c_1}$, i.e., the normalized Type II cost. For the remainder of this chapter, we assume a non-informative prior, i.e., $p_0 = p_1 = 0.5$, and hence regard $\overline{c}_2$ as the normalized Type II cost that is the ratio of Type II cost to Type I cost.

A non-informative prior is consistent with the "equipoise" principle of two-arm clinical trials [53]. However, in some cases we can formulate more informed priors based on information accumulated through earlier-phase trials and other sources. In such cases, the randomization of patients should reflect this information—especially, for life-threatening conditions—for ethical reasons, and the natural framework for doing so is a Bayesian adaptive design [28, 54]. Although this framework is beyond the scope of our current analysis, BDA can easily be applied to adaptive designs and we will consider this case in future research.

The optimal sample size n^* and critical value λ_n^* are determined such that the expected cost of the trial, given by (3.5), is minimized (see Appendix B.1 for a detailed description). The fixed-sample test with these two parameters, i.e., $\texttt{fxd}(n^*, \lambda_n^*)$, will be referred to as the BDA-optimal fixed-sample test. Furthermore, given any fixed-sample test, $\texttt{fxd}(n, \lambda)$—and assuming the test is a BDA-optimal test for a disease with unknown severity (Type II cost) and prevalence—we can impute the severity of disease and its prevalence (see Appendix B.2) implied by the threshold λ.

■ 3.5 Estimating the Cost of Disease

In this section, the two cost parameters, c_1 and c_2, associated with adverse effects of medical treatment and severity of the disease to be treated, respectively, are estimated. To estimate these two parameters, we use the U.S. Burden of Disease Study 2010 [43], which follows the same methodology as of the comprehensive Global Burden of Disease Study 2010 (GBD 2010), however, with only U.S.-level data. Since only the ratio of c_2 over c_1, i.e., $\overline{c}_2$, appears in the expected cost of the trial in (3.5), we use the severity estimates of adverse effects of medical treatment and of disease in the U.S. for c_1 and c_2, respectively.

One of the key factors in quantifying the burden of disease and loss of health due to different diseases and injuries in the GBD 2010 and the U.S. Burden of Disease Study is the YLD (years lived with disability) attributed to each disease in the study population. To compute YLDs, these studies first specify different sequelae (outcomes) for each specific disease, and then multiply the prevalence of each sequela by its disability weight, which is a measures of severity for each sequela and ranges from 0 (no loss of health) to 1

(complete loss of heath, i.e., death). For example, the disability weight associated with mild anemia is 0.005; for the terminal phase of cancers without medication, the weight is 0.519. These disability weights are robust across different countries and different social classes [55], and the granularity of the sequelae is such that the final YLD number for the disease is affected by the current status of available treatments for the disease. This makes YLDs especially suitable for our work, because c_2 is the severity of the disease to be treated, taking into account the current state of available therapies for the disease. We estimate the overall severity of disease using the following equation:

$$c_2 = \frac{D + \text{YLD}}{D + N}, \tag{3.6}$$

where D is the number of deaths caused by the disease, YLD is the number of YLDs attributed to the disease and N is the prevalence of the disease in the U.S., all in 2010. It should be noted that YLDs are computed only from non-fatal sequelae; hence, to quantify the severity of each disease, we add the number of deaths (multiplied by its disability weight, i.e., 1) to the number of YLDs and divide the result by the number of people afflicted with, or who died from, the disease in 2010, hence $D + N$ in the denominator. Furthermore, instead of using the absolute numbers for death, YLD, and prevalence, we use their age-standardized rates (per 100,000) to get a severity estimate that is more representative of the severity of the disease in the population. Age-standardization is a stratified sampling technique, in which different age groups in the population are sampled based on a standard population distribution proposed by the World Health Organization (WHO) [56]. This technique facilitates meaningful comparison of rates for different populations and diseases.

To estimate c_1, which is the current cost of adverse effects of medical treatment per patient, we insert the corresponding numbers for the adverse effect of medical treatment in the U.S. from the U.S. Burden of Disease Study 2010 [43] into (3.6), and the result is $c_1 = 0.07$. It is worth noting that the value of c_1 can be made more precise and tailored to the drug under test if the information from earlier clinical phases, e.g., Phase I and Phase II, is used. However, for simplicity, we only consider a universal value for c_1 for all diseases.

■ 3.6 BDA-Optimal Tests for the Most Deadly Diseases

Using (3.6) and the YLD, death and prevalence rates reported in the U.S. Burden of Disease Study 2010 [43], we can now estimate the severity of some of the leading causes

of death in the U.S. in 2010. Using the estimated severity of each disease, we can then determine the BDA-optimal fixed-sample test for a drug intended to treat that disease. The drug is assumed to have either a positive effect on the disease (corresponding to $\delta_0 = \frac{\sigma}{8}$) or no effect with adverse side effects (corresponding to $\delta = 0$).

The leading causes of death, listed in Table 3.2, are determined in [43] by ranking diseases and injuries based on their associated YLLs (Years of Life Lost due to premature death) in the U.S. in 2010. The following categories, while among the leading causes of premature mortality in the U.S., are omitted from Table 3.2 either because they are not diseases or because they are broad collections (their U.S. YLL ranks are listed in parentheses): road injury (5), self harm (6), interpersonal violence (12), preterm birth complications (14), drug-use disorders (15), other cardiovascular/circulatory diseases (17), congenital anomalies (19), poisonings (26), and falls (29). We have also divided two categories into subcategories in Table 3.2: stroke is listed as ischemic stroke (3a) and non-ischemic stroke (3b), and lower respiratory tract infections is divided into four diseases (11a)–(11d). These choices yield 25 leading causes of death for which we compute BDA-optimal thresholds and compare them to more traditional values.

The estimated severity for each disease, c_2, is reported in the fourth column of Table 3.2. As can be seen, some cancers are not quite as severe as other non-cancerous diseases. For instance, prostate cancer ($c_2 = 0.05$), is much less harmful than cirrhosis ($c_2 = 0.49$), which must be due to the current state of medication for prostate cancer and the lack of any effective treatment for cirrhosis in the U.S. On the other hand, some cancers are shown to be extremely deadly, e.g., pancreatic cancer with $c_2 = 0.71$. Using this measure of severity, we have an objective data-driven framework where different diseases with different afflicted populations can be compared with one another.

Having estimated the severity of different diseases, we apply the methodology introduced in Section 3.4 to determine BDA-optimal fixed-sample tests for testing drugs intended to treat each disease listed in Table 3.2. The sample size, critical value, size, and statistical power of these BDA-optimal tests are reported in Table 3.2. For comparison, we have also listed the imputed prevalence and severity for three conventional 2.5%-level fixed-sample tests in the last three rows of Table 3.2 under the assumption that these conventional thresholds are BDA-optimal (see Appendix B.2).

Some of the diseases listed in Table 3.2 are no longer a single disease but rather a collection of diseases with heterogeneous biological and genetic profiles, and with distinct patient populations [8, 57], e.g., breast cancer. This trend towards finer and finer stratifications is particularly relevant for oncology, where biomarkers have subdivided

Table 3.2. Selected diseases from the 30 leading causes of premature mortality in the U.S., their rank with respect to their U.S. YLLs, prevalence, and severity. The sample size and critical value for the BDA-optimal fixed-sample tests as well as their size and statistical power at the alternative hypothesis are reported.

YLL Rank	Disease Name	Prevalence (Thousands)	Severity	Optimal Sample Size	Optimal Critical Value	Size (%)	Power[a] (%)
1	Ischemic heart disease	8,895.61	0.12	2,028	1.845	3.25	98.36
2	Lung cancer	289.87	0.45	1,373	1.055	14.56	98.68
3a	Ischemic stroke	3,932.33	0.15	1,936	1.744	4.06	98.40
3b	Hemorrhagic/other non-ischemic stroke	949.33	0.16	1,902	1.709	4.37	98.40
4	Chronic obstructive pulmonary disease	32,372.11	0.06	2,343	2.177	1.47	98.22
7	Diabetes	23,694.90	0.05	2,387	2.221	1.32	98.20
8	Cirrhosis of the liver	78.37	0.49	1,300	0.969	16.64	98.67
9	Alzheimer's disease	5,145.03	0.18	1,845	1.640	5.05	98.45
10	Colorectal cancer	798.90	0.15	1,905	1.714	4.33	98.40
11a	Pneumococcal pneumonia	84.14	0.30	1,550	1.311	9.49	98.49
11b	Influenza	119.03	0.20	1,744	1.552	6.03	98.38
11c	H influenzae type B pneumonia	21.15	0.26	1,453	1.279	10.04	98.17
11d	Respiratory syncytial virus pneumonia	14.90	0.07	1,491	1.692	4.53	95.73
13	Breast cancer	3,885.25	0.05	2,374	2.212	1.35	98.19
16	Chronic kidney disease	9,919.02	0.04	2,447	2.283	1.12	98.17
18	Pancreatic cancer	22.67	0.71	1,027	0.587	27.86	98.76
20	Cardiomyopathy	416.31	0.17	1,853	1.659	4.86	98.41
21	Hypertensive heart disease	185.26	0.27	1,633	1.401	8.06	98.50
22	Leukemia	139.75	0.21	1,724	1.522	6.40	98.41
23	HIV/AIDS	1,159.58	0.10	2,087	1.915	2.77	98.31
24	Kidney cancers	328.94	0.12	2,011	1.846	3.24	98.29
25	Non-Hodgkin lymphoma	282.94	0.13	1,944	1.772	3.82	98.32
27	Prostate cancer	3,709.70	0.05	2,414	2.252	1.22	98.17
28	Brain and nervous system cancers	59.76	0.30	1,524	1.290	9.86	98.46
30	Liver cancer	31.27	0.44	1,302	1.004	15.77	98.56
—	2.5%-level Fixed-Sample (85% power)	15.12	0.02	1,150	1.960	2.50	85.02
—	2.5%-level Fixed-Sample (90% power)	17.51	0.02	1,345	1.960	2.50	90.00
—	2.5%-level Fixed-Sample (95% power)	24.60	0.04	1,664	1.960	2.50	95.01

Abbreviations. YLL: Number of years of life lost due to premature mortality, BDA: Bayesian decision analysis.

[a] The alternative hypothesis corresponds to $\delta_0 = \frac{\sigma}{8}$.

certain types of cancer into many subtle but important variations [57]. However, because burden-of-disease data are not yet available for these subdivisions, we use the conventional categories in Table 3.2, i.e., where each cancer type is decided based on the organ host of the tumor.

The reported values for the power of BDA-optimal tests are quite high (all but one have power larger than 98%). This is because the overall burden of disease ($C_2 = Nc_2$) associated with each of these diseases is quite high, due to either severity (large c_2), e.g., pancreatic cancer, or high prevalence (large N), e.g., prostate cancer. Therefore, not approving an effective drug is a costly option by this measurement, hence these BDA-optimal tests exhibit high power to detect positive treatment effects. This general dependence of the statistical power on the overall burden of disease, i.e., its prevalence multiplied by its severity, can be observed in Figure 3.1. In Figure 3.1a, the contour plot of the power of BDA-optimal tests is presented, where most of the contour lines coincide with constant overall burdens of disease, i.e., $Nc_2 = cte$, which are straight lines with negative slope on a log-log graph. Also, to facilitate visualizing where each disease in Table 3.2 lies in the prevalence-severity plane, we have superimposed the YLL rank of each disease in Figure 3.1a. For example, pancreatic cancer is number 18, which has the highest severity among the listed diseases. We have also included the cross-sections of power for BDA-optimal tests in Figures 3.1b and 3.1c.

In sharp contrast to the consistently high power for the BDA-optimal tests in Table 3.2, the size of these tests varies dramatically across different diseases. As is seen in Table 3.2, with few exceptions, the size of the test mainly depends on the severity of the disease. In general, as the severity of the disease increases, the critical value to approve the drug becomes less conservative, i.e., it becomes smaller. This is because the cost per patient of not approving an effective drug becomes much larger than the cost per patient associated with adverse side effects. Consequently, the probability of Type I error, i.e., the size of the test, increases. For example, for pancreatic cancer, the critical value is as low as 0.587, while for the conventional 2.5%-level fixed-sample test it is 1.960. This results in a relatively high size (27.86%) for the BDA-optimal test for a drug intended to treat pancreatic cancer, consistent with the necessity for greater willingness to approve drugs intended to treat life-threatening diseases that have no existing effective treatment.

However, it should be noted that the conventional value of 2.5% for the probability of Type I error, while too conservative for terminal diseases, is not conservative enough for less severe diseases, e.g., diabetes, for which the size of the BDA-optimal test is 1.32%.

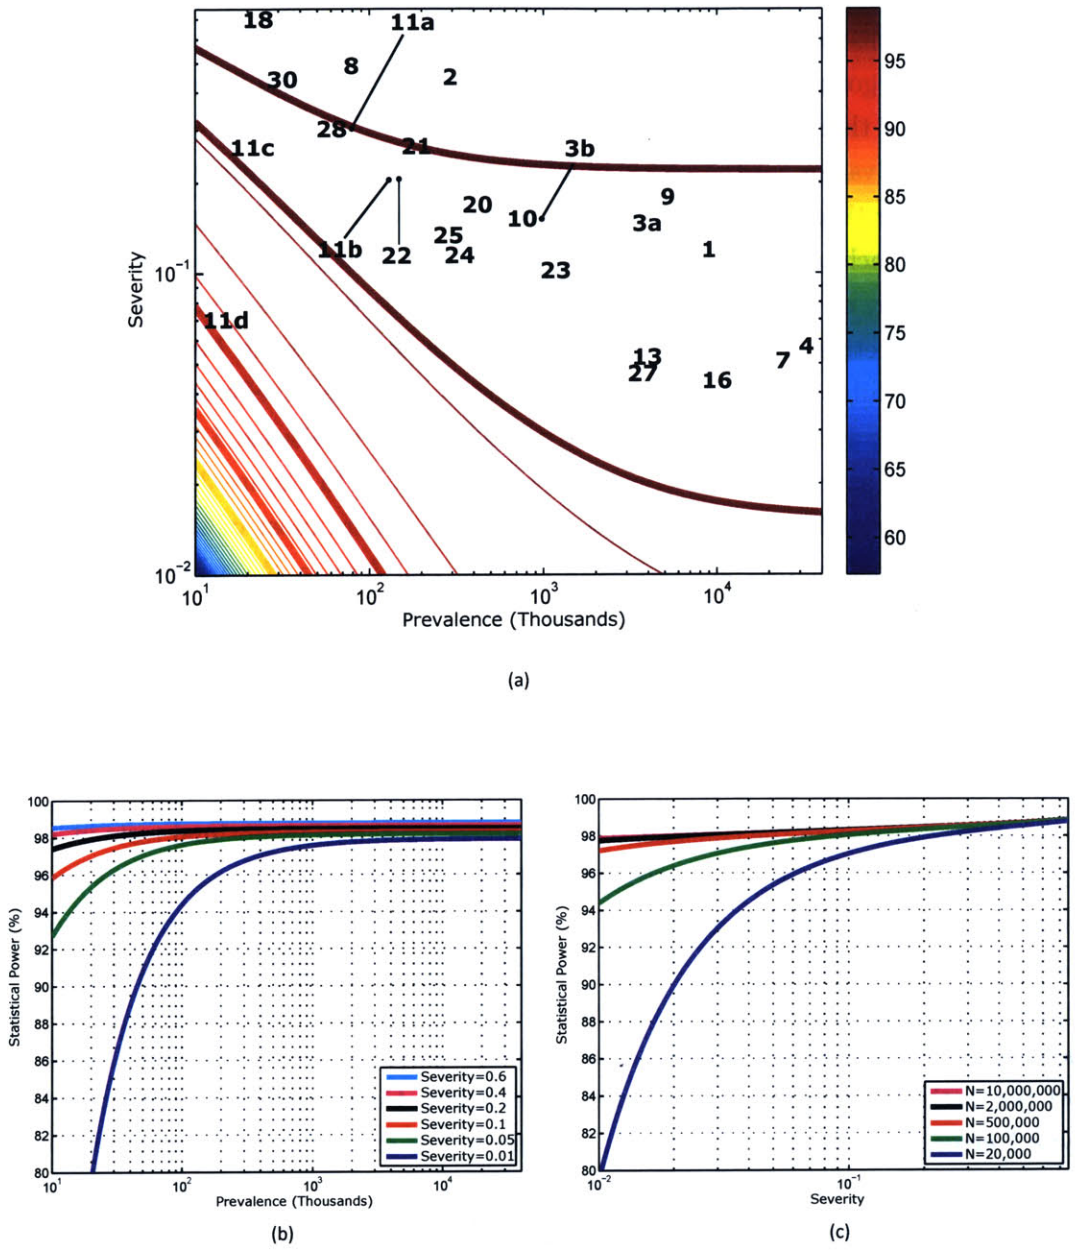

Figure 3.1. The statistical power of the BDA-optimal fixed-sample test at the alternative hypothesis.

Note. Panel (a) shows the contour levels for the power, while panels (b) and (c) demonstrate its cross-sections along the two axes. The contour lines corresponding to the power levels $1 - \beta = 85\%, 90\%, 95\%, 98\%$, and 98.5% are highlighted in panel (a). The superimposed numbers in panel (a) denote the YLL rank of each disease in Table 3.2. The alternative hypothesis corresponds to $\delta_0 = \frac{\sigma}{8}$. YLL: Number of years of life lost due to premature mortality, BDA: Bayesian decision analysis.

The size of BDA-optimal tests for a large range of severity and prevalence values is presented in Figure 3.2. The size monotonically increases with the severity of disease for any given prevalence, and as seen in Figures 3.2a and 3.2b, it becomes independent of the prevalence for all target populations with more than 200,000 patients, hence the horizontal contour lines for x values larger than 200 in Figure 3.2a. This insensitivity of the size to the prevalence of disease makes our model quite robust against estimation noise in the disease prevalence.

It is useful to investigate the dependence of the sample size of BDA-optimal tests on the prevalence and severity of disease. First, we observe in Figure 3.3b that, for any given severity value, the sample size of the BDA-optimal test increases with the prevalence of the disease. This supports the intuitive argument that the sample size should increase with the size of the target population. Furthermore, a unique trend is observed in Figure 3.3c: as the severity of the disease increases, for a large enough target population ($N > 500{,}000$), the optimal sample size continuously shrinks to avoid any delay in getting the effective drug into the market because of the high toll ($C_2 = Nc_2$) that the disease has on society.

On the other hand, for relatively small populations, e.g., $N = 20{,}000$, the optimal sample size peaks somewhere in the middle of the severity spectrum. This occurs because of two opposing trends. The disease burden on society is quite low for small populations and a disease of low severity, hence being exposed to toxic treatment in the trial is not worth the risk. Under these conditions, the sample size should be as small as possible. However, for small populations and a disease of high severity, i.e., a large overall burden of disease, the risk of taking inferior treatment in the trial becomes much smaller than that of waiting for an effective treatment to be approved. Hence, the sample size for $N = 20{,}000$ over very large severity values decreases as severity increases.

In between these two extremes, where the overall burden of disease is not that high, and the disease has intermediate severity, the sample size of the trial is allowed to become larger to guarantee an appropriate balance between approving an effective drug as fast as possible and not exposing the patients to a drug with adverse side effects.

It is worth emphasizing that, as with the size of the test, the sample size of BDA-optimal tests is quite insensitive to the disease prevalence for large target populations (hence, horizontal contour lines in Figure 3.3a over large values of prevalence), which suggests that these results are robust.

Finally, inspecting the conventional fixed-sample tests and the disease prevalence

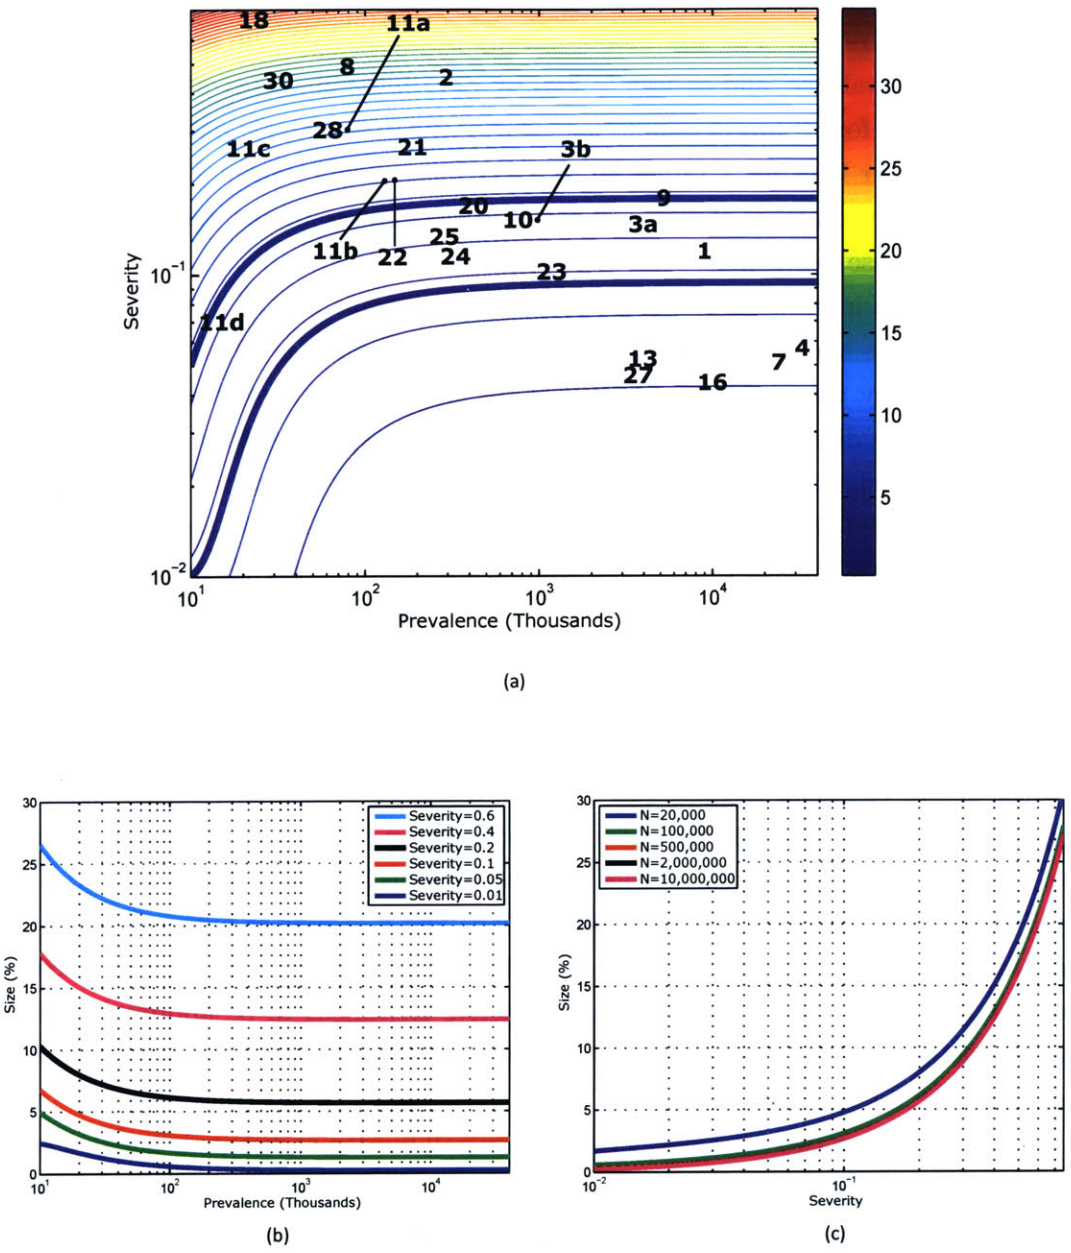

Figure 3.2. The size of the BDA-optimal fixed-sample test as a function of disease severity and prevalence.

Note. Panel (a) shows the contour levels for the size, while panels (b) and (c) demonstrate its cross-sections along the two axes. The contour lines corresponding to $\alpha = 2.5\%$ and $\alpha = 5.0\%$ are highlighted in panel (a). The superimposed numbers in panel (a) denote the YLL rank of each disease in Table 3.2. YLL: Number of years of life lost due to premature mortality, BDA: Bayesian decision analysis.

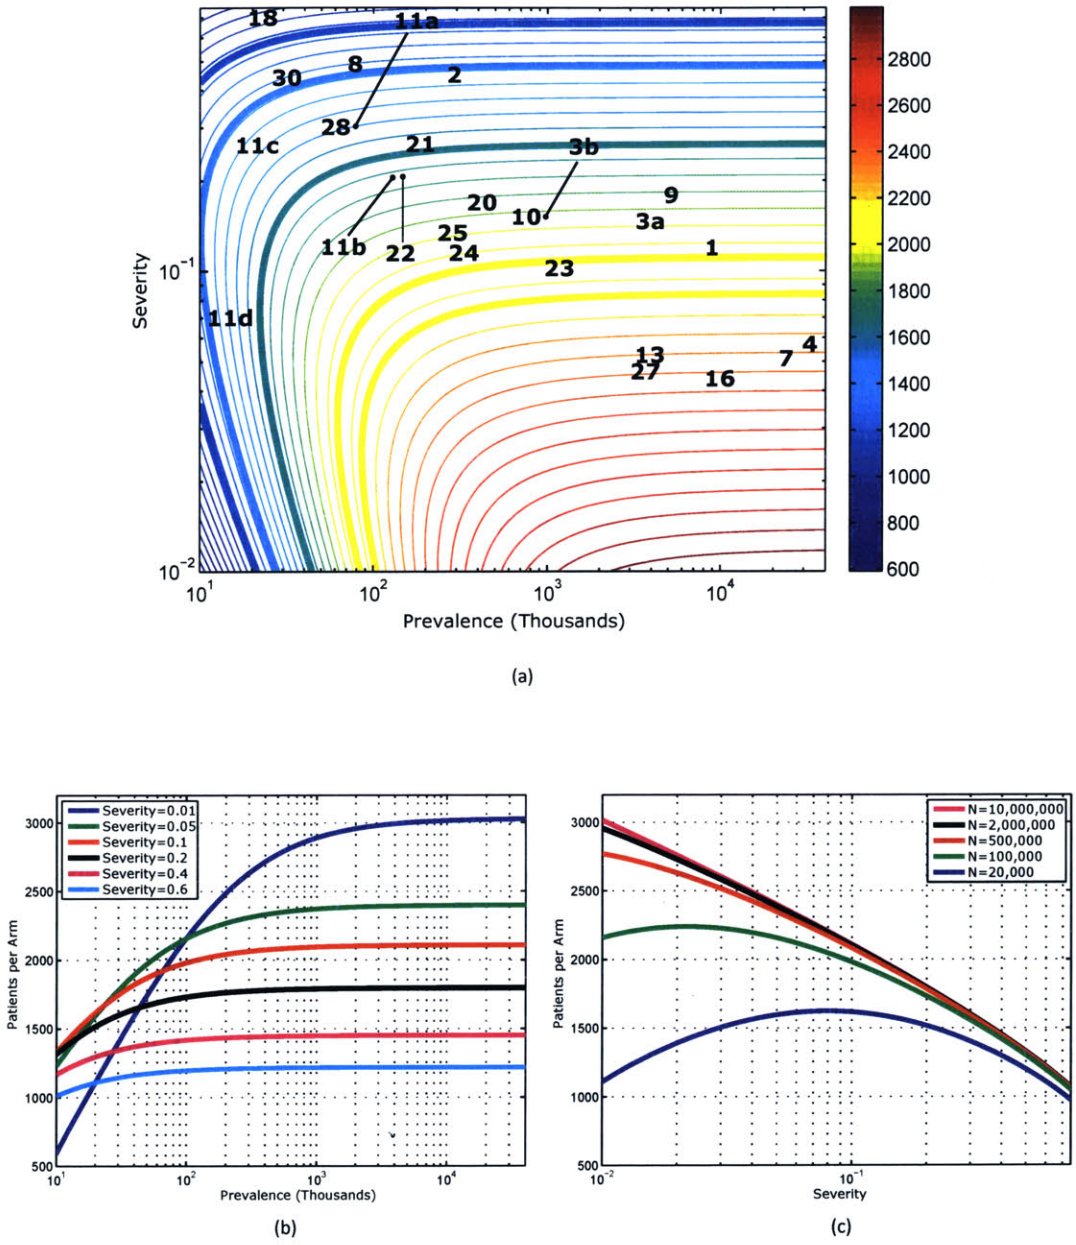

Figure 3.3. The sample size of the BDA-optimal fixed-sample test for different severity and prevalence values.

Note. Panel (a) shows the contour levels for the size, while panels (b) and (c) demonstrate its cross-sections along the two axes. The contour lines associated with the sample size of conventional fixed-sample tests with $\alpha = 2.5\%$ and $1 - \beta = 85\%, 90\%, 95\%, 98\%$, and 98.5% are highlighted in panel (a). The superimposed numbers in panel (a) denote the YLL rank of each disease in Table 3.2. YLL: Number of years of life lost due to premature mortality, BDA: Bayesian decision analysis.

and severity implied by them in Table 3.2 highlights the conservatism of current regulatory requirements imposed on clinical trials and their conduct if we assume that these values are BDA-optimal (see Appendix B.2).

■ 3.7 Conclusion

To address the inflexibility of traditional frequentist designs for clinical trials, we propose an optimal fixed-sample test within a BDA framework that incorporates both the potential asymmetry in the costs of Type I and Type II errors, and the costs of ineffective treatment during and after the trial. Assuming that the current FDA standards represent BDA-optimal tests, the imputed costs implicit in these standards are overly conservative for the most deadly diseases and overly aggressive for the mildest ones. Therefore, changing the one-size-fits-all statistical criteria for FDA drug approval is likely to yield greater benefits to a greater portion of the population.

The BDA framework proposed in this chapter also fills a need mandated by the fifth authorization of the Prescription Drug User Fee Act (PDUFA) for an enhanced quantitative approach to the benefit-risk assessment of new drugs [46]. Due to its quantitative nature, BDA provides transparency, consistency, and repeatability to the review process, which is one of the key objectives in PDUFA. The sensitivity of the final judgment to the underlying assumptions, e.g., cost vs. benefit, can be easily evaluated and made available to the public, which renders the proposed framework even more transparent. However, the ability to incorporate prior information and qualitative judgments about relative costs and benefits preserves important flexibility for regulatory decision-makers.

In fact, a Bayesian approach is ideally suited for weighing and incorporating patient perspectives into the drug-approval process. The 2012 Food and Drug Administration Safety and Innovation Act (FDASIA) [58] has "recognized the value of patient input to the entire drug development enterprise, including FDA review and decision-making." One proposal for implementing this aspect of FDASIA is for the FDA to create a patient advisory board consisting of representatives from patient advocacy groups, with the specific charge of formulating explicit cost estimates of Type I and Type II errors. These estimates can then be incorporated into the FDA decision-making process, not mechanically, but as an additional inputs into the FDA's quantitative and qualitative deliberations.

To incorporate other perspectives from the entire biomedical ecosystem, the membership of this advisory board could be expanded to include representatives from other

stakeholder groups—caregivers, physicians, biopharma executives, regulators, and policymakers. With such expanded composition, this advisory board could play an even broader role than the concept of a Citizens Council adopted by NICE.[3] The diverse set of stakeholders can provide crucial input to the FDA/EMA, reflecting the overall view of society on critical cost parameters. However, the role of such a committee should be limited to advice; drug-approval decisions should be made solely by FDA officials. The separation of recommendations and final decisions helps ensure that the adaptive nature of the proposed framework will not be exploited or gamed by any one party.

In fact, because of its role as the trusted intermediary in evaluating and approving drug applications, the FDA is privy to information about current industry activity and technology that no other party possesses. Therefore, the FDA is in the unique role of formulating highly informed priors on various therapeutic targets, mechanisms, and R&D agendas. Applying such priors in the BDA framework could yield very different outcomes from the uniform priors we used in Section 3.6, which assumes a 50/50 chance that a drug candidate is effective. While 50/50 may seem more equitable, from a social welfare perspective it is highly inefficient, potentially allowing many more expensive clinical trials to be conducted than necessary. Although the FDA cannot be expected to play the role of social planner, and should be industry neutral in its review process, nevertheless, ignoring scientific information in favor of 50/50 does not necessarily serve any stakeholder's interest. Moreover, using 50/50 when more informative priors are available could be considered unethical in cases involving therapies for terminal illnesses. For example, for pancreatic cancer, if the prior probability of efficacy is 60% instead of 50%, the size of the BDA-optimal test would be 51.2% rather than 27.9%, leading to many more approvals of such therapies. The BDA framework can yield decisions that are both more economically efficient and more humane.

Finally,the drug-approval process is not always a binary choice, and in such cases, the BDA framework can be extended by defining costs for a finer set of events. In fact, the variability of drug response in patient populations—attributed to biological and behavioral factors—has been recognized as a critical element in causing uncertainty and creating the so-called "efficacy-effectiveness" gap [59] (where efficacy refers to therapeutic performance in a clinical trial and effectiveness refers to performance in practice). Several proposals have been made for integrated clinical-trial pathways to bridge this gap [60].

Moreover, new paradigms have also been proposed to address the risk associated

[3]See `https://www.nice.org.uk/Get-Involved/Citizens-Council`.

with the binary nature of the current approval process, e.g., staggered approval [61, 62] and adaptive licensing [63], which the EMA is actively pursuing [64]. In fact, one of the design principles called for by [63] is less stringent statistical significance levels to be employed in efficacy trials for drugs targeting life-threatening diseases and/or rare conditions. Our BDA framework provides an explicit quantitative method for implementing this principle. The fact that the adaptive pathway has great potential to benefit all key stakeholders [65] provides more motivation for employing BDA in the drug-approval process.

Chapter 4

Reimbursement for Curative Therapies

There has been much debate among several stakeholder communities over the prices of certain drugs, especially those that treat serious conditions such as cancer, Hepatitis C Virus (HCV) infection, and rare diseases. Apart from important ethical and social issues surrounding the affordability of life-saving therapies, there is a more practical issue of how payments can be financed. We propose to address the financing issue through healthcare loans (HCLs), which are the equivalent of mortgages for large healthcare expenses such as the cost of curing a disease using highly personalized therapies that are not fully covered by insurance. We then propose using securitization—an efficient financial engineering method— to finance a large pool of these HCLs using both debt and equity. Using numerical simulations, we demonstrate that the proposed financing vehicle is viable under a wide array of economic environments and cost parameters, making new therapies accessible to a much broader patient population. Moreover, the proposed framework better aligns the interests of all stakeholders, and has the potential to accelerate biomedical innovation while reducing aggregate healthcare costs.

■ 4.1 Introduction

Much has been written about the failings of our current drug development pipelines. Many new drugs offer little or no benefit over currently available and less expensive alternatives. Others, most commonly among cancer therapies, offer very small improvements in outcome at high prices. In contrast with these incremental advances, new therapies that truly transform medical care (e.g., penicillin for pneumococcal infection, highly active anti-retroviral therapy for human immunodeficiency virus (HIV), imatinib for chronic myelogenous leukemia (CML)) are few and far between. When

these transformative products emerge, there is a compelling urgency to extend them as broadly as possible. Doing so introduces a new challenge, specifically the cost of offering treatment to all who could benefit.

Two recent advances have attracted significant attention from the medical and lay communities. The first is curative therapy with antiviral drugs that target NS3/4A protease, NS5B polymerase and NS5A replication complex in patients with Hepatitis C Virus (HCV). Over 90% of infected individuals are believed to be cured of HCV infection with only 6-8 weeks of combinations that include these agents at a price of around $84,000 [66]. Extending curative treatment to all 2 million Americans with chronic HCV infection at that price would cost $168 billion. Extending curative therapy to all 180 million people worldwide with HCV infection would cost over $15 trillion.

The second example of a truly transformative therapy is chimeric antigen receptor-T (CAR-T) cells, which may cure up to 85% of children with relapsed and refractory acute lymphoblastic leukemia (ALL) after a single infusion. These children have no other options and would otherwise die of their disease. The price of a CAR-T cell infusion has not been set as the Food and Drug Administration (FDA) has not yet licensed any CAR-T cell products. Approvals are likely to be forthcoming, so it will soon be necessary to ask, what is an appropriate price for a single infusion that extends the life of a 3-year old child by 80 years? On the basis of quality-adjusted life years (QALYs), using $50,000/QALY as an accepted standard, the back-of-envelope calculation argues that $4 million may be appropriate. As a benchmark, the fund to compensate families of those killed on 9/11 paid an average of $2 million per victim. This provides further support that our society values a single American life at amounts in this range.

If one accepts that the imperative is to offer truly transformative therapies, as quickly and broadly as possible, creative solutions are clearly needed. This is true whether the therapy is directed against an infectious disease with a large prevalence that is also a public health concern (like HCV or HIV), or an orphan disease that is invariably fatal (like refractory ALL or CML). The Lancet Commission on Addressing Liver Disease in the U.K. recently estimated that with new therapies, we could "contemplate the eradication of infections from chronic hepatitis C virus in the U.K. by 2030." This is a tremendously under-ambitious goal, considering that it would permit the morbidity, mortality, and spread of HCV to continue for 15 years when the only factor preventing eradication on a much shorter timeline is the aggregate upfront cost. Insurance companies would not have too big a problem providing access to these therapies for future patients because the incidence rate of the disease can be quite precisely

estimated for coming years, and the insurance premiums can be adjusted accordingly. However, covering the existing patients, the so-called "warehoused" patients, is highly burdensome for insurance companies right after a new therapy emerges in the market. This is the main reason that some patients are currently denied coverage for their therapies—they are not "sick enough."

To alleviate the issue of large aggregate upfront payments for insurance companies, we propose healthcare loans (HCLs)—the equivalent of drug mortgages—for patients to cover the portion of the drug's price that insurance companies cannot pay for. Like home mortgages, HCLs are loans taken out by patients that they repay over a period of time following their treatment. We then use portfolio theory and financial engineering techniques [13, 67] to design an efficient financing vehicle for these healthcare loans and to reduce the borrowing cost for the patients to lowest possible levels. Insurance companies, in addition to paying a portion of the therapy's price upfront, would play a critical role in ensuring that the proposed vehicle functions as planned so as to attract the cheapest type of financing, i.e., highly-rated debt.

There have previously been other proposals on payers making annuity-based payments for curative therapies, specifically gene therapies [68]; however, to the best of our knowledge, we are first to propose a detailed and practical annuity-based payment structure, where patients rather than payers make periodic payments.

■ 4.2 Portfolio Theory

Suppose a $40,000 loan is granted to a patient, and he/she is required to repay it by making annual payments of $6,700 for 9 consecutive years following the loan origination, referred to as the repayment period. Because during the repayment period, the borrower might stop their payments—due to either default on the loan or death—there are distinct scenarios for the stream of payments that can occur. For example, the lender might receive one payment of $6,700 in the first year, and nothing afterwards. In another scenario, the lender may receive all the payments over the repayment period as scheduled. For simplicity, we assume that, in each year, the scheduled payment is either fully made or missed, and if a payment is missed, all subsequent payments will be missed too. Any investment that yields different streams of cash flows under different conditions is referred to as a "risky" investment [69]. The natural question to ask is: what is the value of this future risky stream of payments from the lender's perspective? Clearly, if the lender values this investment below $40,000, he/she will not lend under

these circumstances.

In financial economics, it is common to use the discounted cash flow valuation (net present value) methodology to value multi-period risky investments including the example considered here [69, 70]. This approach can be divided into two separate parts: First, the net present value of the cash flow stream associated with each scenario is derived by properly discounting future cash flows (net of costs) to present (because \$1 next year is less valuable than \$1 today). For example, the present value of the cash flow stream in the case where the borrower stops their payments after making the first k payments, denoted by PV_k, is given by:

$$PV_k = \sum_{i=1}^{k} \frac{X}{(1+R)^i} = \frac{X}{R}\left[1 - \left(\frac{1}{1+R}\right)^k\right], \quad k = 0, 1 \ldots, n, \tag{4.1}$$

where X = \$6,700 and $n = 9$ years, and $R > 0$ represents the rate of return expected (required) by the lender to justify bearing the risk of the investment, and is explained in more detail below. The value of the investment, denoted by V, is then simply the expectation of the present values over all possible scenarios:

$$V \equiv E[PV] = \sum_{k=0}^{n} p_k PV_k = \sum_{k=0}^{n} p_k \frac{X}{R}\left[1 - \left(\frac{1}{1+R}\right)^k\right], \tag{4.2}$$

where $E[\cdot]$ is a shorthand for expectation, and p_k represents the probability that only the first k payments are made and the remaining payments are missed. We derive the final equation by substituting the expression given in (4.1) for PV_k.

Now, suppose that the combined probability of death and default for each year is denoted by $p_0 = 5.0\%$, and defaults/deaths in different years are independent. Then, the probabilities, p_ks, in (4.2) can be derived as the following:

$$p_k = \begin{cases} p_0 (1-p_0)^k, & k = 0, 1, \ldots, n-1 \\ \\ (1-p_0)^k, & k = n \end{cases}. \tag{4.3}$$

Using the expressions for probabilities in (4.3), the value of the investment in (4.2) can be calculated for different expected returns, as presented in Figure 4.1a. In Figure 4.1a, the expected return for which the value is equal to the initial investment (\$40,000) is $R = 3.69\%$.

As observed in Figure 4.1a, the value of the lending, given by (4.2), decreases with the expected return of the lender, R. The higher the expected return, R, the lower the value of the investment, V, from the investor's perspective *ceteris paribus*. The expected rate of return, sometimes referred to as the cost of capital, of the lender represents the annual profit on every dollar that the investor demands to gain, on average, from the initial investment. For example, if $R = 10\%$, the lender requires to obtain 10 cents per year on every invested dollar. If the investment meets this hurdle, i.e., if V in (4.2) is larger than the initial investment ($40,000 in our example), he/she makes the investment; otherwise, making the investment is not justified. Hence, in our example, the lender lends only if $R \leq 3.69\%$.

Different classes of investors have different expectations and risk appetites. And, the larger the investor's risk appetite, the larger the return they demand on the investment. The same holds true for investments: The higher the risk of an investment, the larger the expected return it should offer to be viable. At one extreme, debt investors ask for quite low expected returns (e.g., $R = 2\% - 3\%$), and in exchange, they have a small risk tolerance, i.e., they favor "safe" investments that have fewest cash flow scenarios most, and avoid investments with a lot of possible ups and downs. For example, Treasury bonds issued by the U.S. government are considered as one of the safest investments—because of the full faith and credit in the U.S. government—and they consequently offer a rather small return ($R = 1.51\%$ for 5-year Treasury bonds, updated September 2, 2015). At the other extreme, there are equity investors who are willing to take a lot of risk (a lot of possible downsides and upsides) if the possible upside and consequently the expected return are large enough to justify the risk-taking (e.g., double digit returns, $R \geq 12\%$).

The standard deviation of the cash flows that an investment generates can serve as a proxy for its risk. The solid line and the blue bars in Figure 4.1b demonstrate the expected value and standard deviation of the payments, respectively, for a single loan during the repayment period. To analyze the risk of the investment from another perspective, the blue staircase curve in Figure 4.1c presents the cumulative probability distribution (CDF) of the normalized present value of the stream of payments for a single loan, given by (4.1). The present value is calculated using $R = 3.69\%$, and is normalized by the initial investment, i.e., $40,000. This expected return is used because, as seen in Figure 4.1a, it makes the value of the investment, which is the expectation of the present values, equal to the initial investment. However, while the mean of the blue distribution in Figure 4.1c is 1, there is a significant likelihood of receiving only

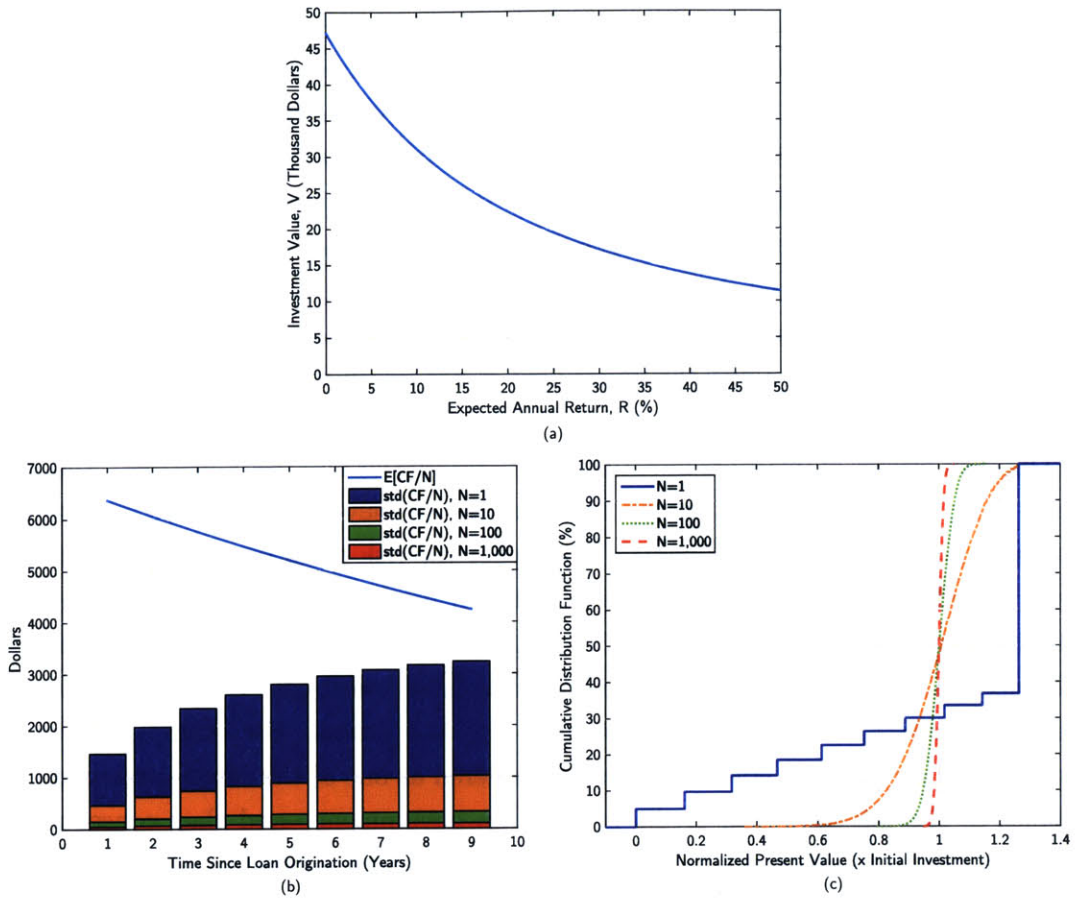

Figure 4.1. (a) Investment value, given by (4.2), over a range of expected returns, R. (b) Expected value and standard deviation of cash flows normalized by the number of loans over the 9 years following the origination of loan(s). (c) The cumulative probability distribution (CDF) of the present value of the investment normalized by the initial investment for a few portfolio sizes, N.

Note. The expressions in (4.2) and (4.3) along with $R = 3.69\%$ are used to calculate present values. Expected value is denoted by $E[\cdot]$, and $std(\cdot)$ represents standard deviation.

a fraction of the initial investment. Visually, the blue curve has a wide spread around the vertical line where the normalized present value is equal to 1.

Therefore, by inspecting Figures 4.1b and 4.1c, lending a single loan is deemed too risky for investors with low expected returns. On the other hand, if investors with higher expected returns who have a larger risk appetite were considered, the risk of the investment might be acceptable for them while its expected value might seem too weak as seen in Figure 4.1a.

The solution to this conundrum is diversification. Consider making loans to N

customers rather than a single borrower, i.e., investing in a portfolio of N loans. The initial investment required is N times larger than for a single loan and so is the expected cash flow in each year. Hence, if normalized by N, the expectation of the cash flow is independent of the size of the portfolio. However, assuming each borrower's payments are to some extent independent of others', the standard deviation of future cash flows does not increase linearly in N. Hence, if normalized by N, the cash received in each year becomes increasingly more certain with the number of loans in the portfolio. This is shown in Figure 4.1b, where each bar graph demonstrates the standard deviation of the normalized cash flow for each portfolio size over the repayment period. The expectation of the normalized cash flows, however, is independent of the portfolio size. The risk reduction can further be seen in Figure 4.1c, where the distribution of the normalized present value becomes increasingly narrower as the portfolio grows in size, while the mean of the distribution remains intact at 1. In financial economics, diversification plays a significant role in risk reduction, and is the centerpiece of the portfolio theory [71, 72]. By reducing the risk of the portfolio through diversification, it is possible to make it attractive to more risk-averse investors such as debt investors, who are the low end of the expected-return spectrum. By appealing to these investors and considering the current near-zero interest rate environment, we can keep the cost of borrowing for the patients (borrowers) at lowest possible levels to control the financial burden of our proposal. In the following section, we explain how to use diversification and to design a "safe" investment for debt investors, who have a small risk appetite, so as to take of advantage of the low cost of debt financing.

■ 4.3 Securitization

As discussed in the previous section, there are different classes of investors, each of which possessing a different risk-return profile. The financing cost can be reduced by shifting to the most risk-intolerant investors, who demand a great deal of protection against risk. The other end of the financing cost spectrum has the greatest risk tolerance while expects a large upside in return, increasing the cost of financing. Securitization is a commonly used financial engineering method that allows us to benefit the most from the whole investor spectrum by issuing different types of debt and equity [67, 73]. Securitization effectively slices the investment into sub-investments, each of which is tailored to a different risk-return profile and is supported by a portion of the assets in the portfolio. By tailoring these sub-investments to different investors' expectations,

securitization allows us to tap into global capital markets to share risk among an enormous pool of investors. By using debt financing, the financing costs can be reduced to historically low levels while using equity and other credit enhancements protects debt investors against capital loss to the extent demanded by the debt investors.

For example, as seen in 4.1c, for large portfolio sizes, there is a 50% chance of having a present value less than the initial investment, regardless of the size of the portfolio. Therefore, if the investors expect a 1% chance of receiving less than promised, the diversification per se does not resolve that issue. Instead, we can raise a portion, e.g., 80%, of the initial investment from these risk-averse investors and the remaining 20% can be raised from equity investors who have a larger risk appetite. As observed in 4.1c, the likelihood of receiving less than 80% of the initial investment decreases with the size of the portfolio, and for large enough portfolios, this likelihood can definitely be less than 1%. To summarize, by using diversification and securitization, we can reduce the cost of financing without becoming overly risk-averse; i.e., we can find a sweet spot in the risk-return spectrum that most closely matches the characteristics of the investment.

■ 4.4 HCL Fund for HCV Curative Therapies as an Illustrative Example

Suppose a drug's price is $84,000 and the insurance company can pay only $44,000 for each drug. In our proposed framework, the insurance company pays this portion of the price and a "special purpose entity" (SPE) is set up to pay the remaining amount ($40,000 per drug) on behalf of the patient so that he/she can get access to the drug immediately (Figure 4.2a). The amount paid by the SPE is effectively a loan (mortgage) granted to the patient that has to be repaid over a pre-specified period of time, referred to as the repayment period (9 years in our simulations). If we assume a 9.1% annual interest rate on the loan, the HCL payments will be $6,700 per year for each borrower (Figure 4.2b). We assume that each patient fulfills their payment obligations until they either die or default on their loan. Under either condition, the payments will stop completely; i.e., the recovery rate on loans with stopped payments is zero. In the following sections, we explain our assumptions with regard to financing, default characteristics, and mortality rates of the patients after going through the course of therapy.

■ 4.4.1 Portfolio Dynamics

The SPE holds a portfolio of claims on the future loan-payment streams (12,500 loans in our simulations) to benefit from diversification as explained in Section 4.2. This is similar to the securitization of home mortgages, student loans, auto loans, and credit card receivables, to name a few. The size of the fund managed by the SPE is \$500mm, where mm denotes million (12,500 × \$40,000 = \$500mm). The SPE usually sets up an independent entity, usually referred to as the fund's general partner, to manage the portfolio; however, for simplicity, we refer to both the manager and this special purpose entity as the SPE. The SPE raises the required \$500mm by selling notes to both equity and bond investors. More precisely, 80% of the fund, i.e., \$400mm, is raised through the sale of senior bonds, which are highly-rated bonds that have low default probabilities and small default losses. These bonds have an annual coupon rate of 2.1% and pay the investors once a year (in the U.S., bonds usually have semi-annual payments, however, for simplicity, we assume only annual payments). Of the remaining 20% of the fund, half (i.e., 10% or equivalently, \$50mm) is raised by selling subordinate (junior) notes that have less protection against possible losses and offer a higher coupon rate in return compared to the senior bonds (e.g., 2.50% in our simulations). In addition to interest payments, both senior and junior bonds repay their principals in nine equal annual installments starting in the first year; i.e., they amortize in the first year. Equity investors invest in the remaining 10% of the fund that amounts to a \$50mm investment. There are not any scheduled payments for these investors; however, they receive cash distributions in each year provided that the fund has excess cash after making scheduled bond principal and interest payments and saving \$1mm in a cash reserve account to preserve the fund's liquidity.

In the governing documents of the fund, the cash flow waterfall specifies the order with which different investors (senior bondholders, junior bondholders, equity investors) receive payments in both normal times and default events. The senior bondholders have the highest priority, followed by junior bondholders, and the equity investors have the lowest payment priority (Figure 4.2b). Therefore, equity investors are first to absorb any losses to the portfolio—due to loan defaults—then junior bondholders, and senior debt investors are last to experience any losses (losses propagate through the capital structure of the fund from the bottom to the top in Figure 4.2b). This provides senior bonds with a degree of protection because they only constitute 80% of the capital employed in the fund, hence, even if the portfolio experiences a loss, so long as the loss is less than 20% of the fund's capital, senior bonds do not experience any default or

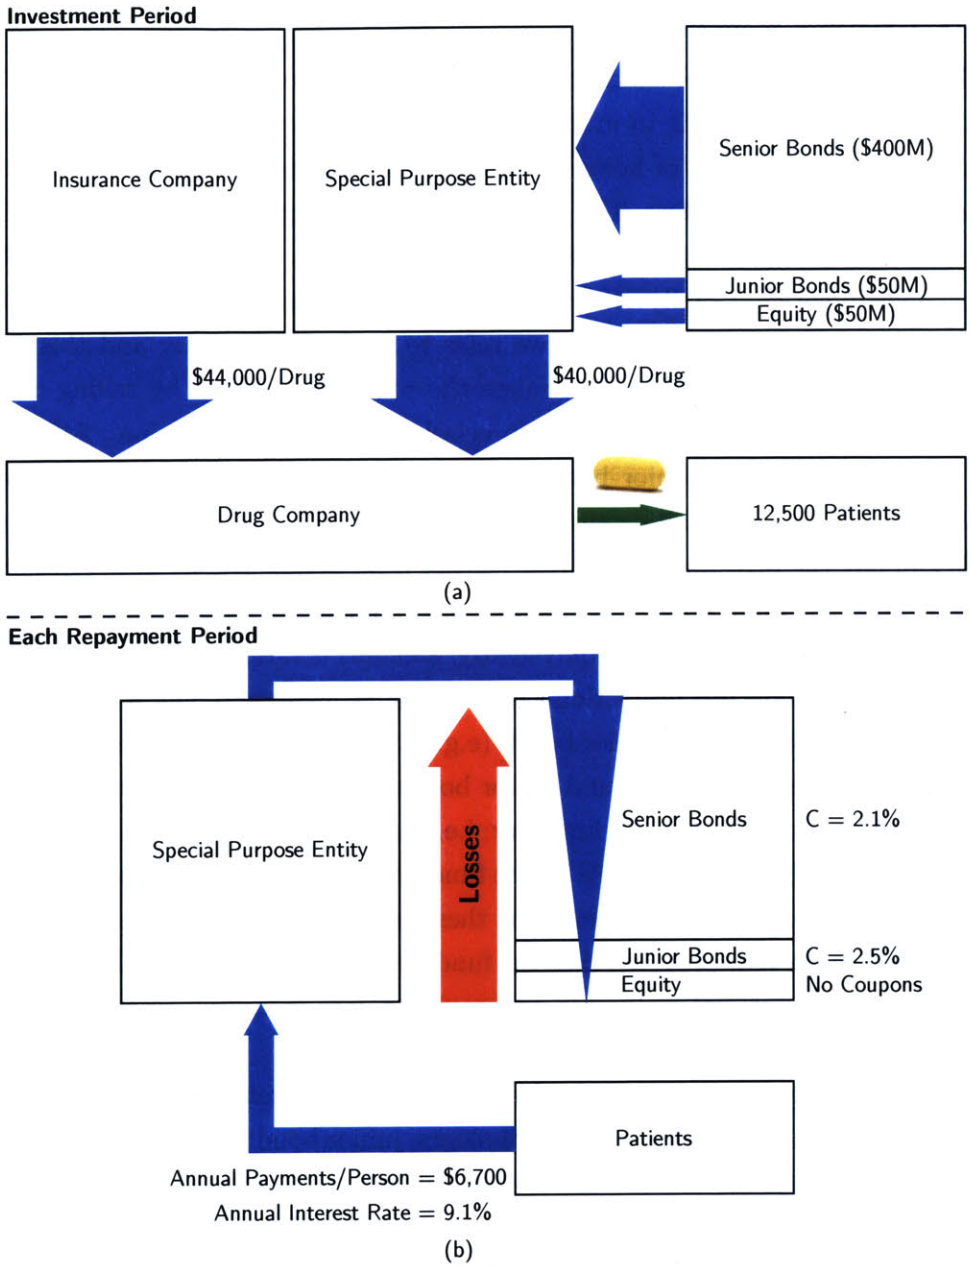

Figure 4.2. Schematic cash flow diagram for the proposed HCL fund. Panel (a) demonstrates the investment period in which the investors buy the notes issued by the special purpose entity (SPE), and using the cash raised from the sale of the notes, the SPE pays a portion of the drug's price, and the patients receive the curative therapy. The bottom panel (b) shows the flow of cash in each repayment period, in which the patients make their annual loan payments, and the investors receive cash payments based on the seniority of their notes. The losses propagate from the bottom to the top.

Note. The source of the pill picture is Hepatitis C Online (`http://www.hepatitisc.uw.edu`).

loss. This layering of the capital structure is usually referred to as subordination, and is commonly used to protect the most senior bonds. In addition to the subordination employed in the fund's capital structure, we use extra credit enhancement techniques, such as interest coverage and overcollateralization tests as well as cash flow diversion mechanisms, to provide additional protection to the bonds, especially the senior bonds. For more detail on these techniques, see Appendix C.1.

As the last layer of protection for bondholders, we assume a third-party guarantor for the bonds issued by the SPE, promising to make scheduled interest and principal payments were the SPE unable to do so. Having guarantees on the bonds helps insure debt investors against extreme events, and consequently, the bonds' coupon rates and therefore financing costs can be reduced to decrease the financial burden of the loans on the patients (for a detailed description of the relation between risk and expected return see Section 4.2).

■ 4.4.2 Loan Default Dynamics

Considering that the debt-payment-to-income ratio, defined as the ratio of annual debt payments to annual income, is a commonly used proxy for debt burden [74], we assume that the borrower's expected default probability increases with this ratio. We design a statistical model for the relationship between a loan's default probability and the debt-payment-to-income ratio attributed to that loan, and we then calibrate the model parameters using federal student loan data as explained in Appendix C.2. Because different borrowers have possibly different incomes, the dependence of the expected default probability on the debt-payment-to-income ratio creates a realistic heterogeneity among loan defaults in the portfolio, which is crucial for any statistical modeling of a portfolio of loans. In addition, we estimate the probability distribution of the annual household income for patients afflicted with HCV using the Chronic Hepatitis Cohort Study (CHeCS) data on U.S. patients with chronic HCV infection [75]. Using the estimated income distribution, the annual loan payments, and the relationship between expected probability of loan default and borrower's debt-payment-to-income ratio, we can then model the credit risk of the portfolio of loans in our proposed framework. We further make the loan defaults in the portfolio correlated by making all the expected loan default probabilities in each year rise or fall together within a pre-specified range (75% to 125% of the initially assumed values). The induced default correlation is larger for lower-income patients than the patients with a higher household income (e.g., 10% for two patients with the annual household income of $40,000, and 5% for annual incomes of

$80,000), consistent with the intuition that the higher the debt burden—associated with lower incomes—the more correlated the loan defaults due to any changes in the economy and wages. As is known in modern portfolio theory, correlation among loan defaults limits the ability of diversification in reducing the risk of the portfolio indefinitely since it makes the defaults more likely to occur simultaneously rather than independently [72].

Because the debt burden and the expected default probability become larger as the income falls, we set a cutoff income threshold ($35,000) in our simulations. In other words, we assume that all the patients, who are granted an HCL, have annual household incomes of at least $35,000. This is a practical assumption because it is most likely that the patients with household incomes below $35,000 have Medicaid coverage. Our focus here is on private insurance companies only, while a different structure can be used solely for Medicaid patients who would most likely be heavily burdened by the proposed payments in this chapter. In that Medicaid-specific financing vehicle, the U.S. government guarantees the bonds so that they can have the lowest possible coupon rates. On the other hand, the loan repayment period can be made longer than assumed here to bring the payment burden down to an acceptable range. Writing the guarantees would in turn cost the government a small fraction of what Medicaid would have to pay upfront in the current pricing environment (see Table 4.1).

Lastly, to evaluate the performance of the HCL fund under a few conditions, we consider three scenarios for the dependence of the expected probability of HCL default on the borrower's income as presented in Figure 4.3a, and referred to as the pessimistic, baseline, and optimistic scenarios. The baseline case is our original estimated relationship between the probability of default and the patient's income, while the pessimistic and optimistic scenarios assume higher and lower default probabilities, respectively, across almost all the incomes compared to the base case. For a detailed description of these three scenarios, see Appendix C.2.

■ 4.4.3 Post-Medication Mortality Dynamics

There have been many longitudinal studies on post-medication mortality rates for HCV treatments with interferon and/or ribavirin regimens (see [76–78] and references therein). However, because NS5B polymerase inhibitors such as Sovaldi are relatively new, there has not yet been any published study on mortality rates for patients treated with these regimens. On the other hand, the studies on interferon and ribavirin treatments unanimously demonstrate that all-cause mortality rates for patients with a sus-

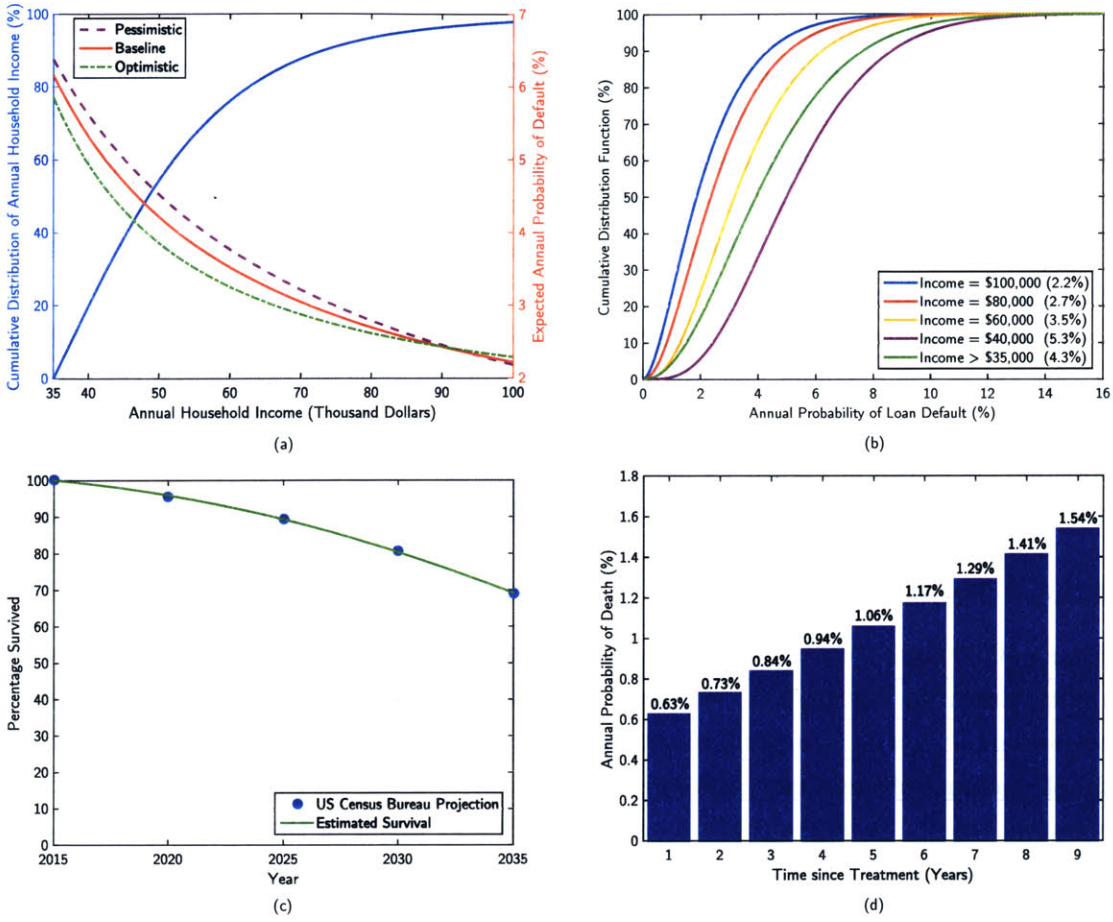

Figure 4.3. Summary of assumptions for loan defaults and post-medication mortality used in the simulations. (a) The cumulative distribution function (CDF) of the annual household income for patients with chronic HCV (left axis), and the estimated expected default probability as a function of income for three different scenarios (right axis). (b) The CDF of annual default probability, in the baseline scenario, for the patient population (income > $35,000) as well as for a few incomes. The numbers in parentheses denote the expected default probability associated with that category. (c) The U.S. Census Bureaus projected numbers for the baby-boomer generation as well as our estimated post-medication survival curve for each patient. (d) The implied annual death probabilities by the survival curve in (c) for the 9 years following the treatment.

tained virological response (SVR) are statistically indifferent from mortality rates of a matched general population [76]. Therefore, considering that more than 90% of the patients treated with the new generation of direct-acting agents achieve a 12-week SVR (SVR12) with less adverse effects than previous standards of care, as shown in the NUETRINO trial [66], we use general population mortality rates as a proxy for patient mortality rates after receiving this new line of treatment [79, 80].

Furthermore, because more than 75% of the HCV-infected population is estimated to be baby-boomers (born between 1945 and 1965),[1] we use the projected population numbers for the U.S. baby-boomer cohort, published by the U.S. Census Bureau [81], to estimate a post-medication survival curve (for a detailed description, see Appendix C.3). The estimated survival curve and annual death probabilities are drawn in Figures 4.3c and 4.3d. Our estimated 10-year survival rate (89.3%) is close to the rates reported elsewhere [76, 80, 82].

To demonstrate the overall impact of HCL defaults and post-medication mortality of patients on the portfolio, the cumulative losses of the portfolio throughout the fund's life for the baseline scenario are presented in Figure 4.4. The cumulative losses of the portfolio are drawn in Figure 4.4 both as a percentage of the number of initial HCLs in the portfolio and as a percentage of the original balance of the HCLs in the portfolio.

■ 4.4.4 Performance Results

We use 10,000,000 Monte Carlo simulation paths per each scenario of HCL default probability, drawn in Figure 4.3a, to evaluate the performance of the HCL fund. Different performance metrics for each scenario are reported in Table 4.1. Total cash return is a simple measure of the equity investment performance, which measures how much equity investors receive throughout the fund's life per every dollar of their initial investment (with no time-discounting), and is defined as the following:

$$\text{TCR} \equiv \frac{\text{Cumulative cash returned to equity investros over the fund's life}}{\text{Initial equity investment}}. \tag{4.4}$$

Another metric for equity performance is the internal rate of return (IRR), which is the solution to the following equation and can serve as a proxy for the annual growth of the equity investment:

$$I_0 = \sum_{k=1}^{n} \frac{c_k}{(1+\text{IRR})^k}, \tag{4.5}$$

where I_0 represents the initial equity investment ($I_0 = \$50\text{mm}$ in our simulations), the life of the fund is given by n (in our simulations, $n = 9$ years) and c_k denotes the cash received by equity investors in year k.

It is clear in Table 4.1 that the equity performance is worst in the pessimistic scenario (IRR = 9.4%, TCR = 131.0%), while the optimistic scenario yields the best equity

[1] See `http://www.cdc.gov/knowmorehepatitis/Media/PDFs/FactSheet-Boomers.pdf`.

Table 4.1. Performance results of the proposed HCL fund for three scenarios of loan defaults.

	Pessimistic[a,b]	Baseline[a,c]	Optimistic[a,d]
Loan default assumption			
Expected annual loan PD (%)	4.6	4.3	4.0
Equity tranche performance			
Expected total cash return (%)	131.0	144.8	161.6
Expected annualized IRR (%)	9.4	12.5	15.8
Pr(IRR $< 0\%$) (%)	1.1	< 0.1	0.0
Pr(IRR $\geq 5\%$) (%)	87.3	98.7	100.0
Pr(IRR $\geq 10\%$) (%)	46.3	79.6	98.4
Pr(IRR $\geq 15\%$) (%)	5.2	22.5	63.6
Pr(IRR $\geq 20\%$) (%)	< 0.1	0.2	4.0
Debt tranches performance			
Senior bonds			
Weighted average life[e] (years)	4.98	4.99	5.00
Probability of default (bps)	< 0.1	< 0.1	< 0.1
Expected loss (bps)	< 0.1	< 0.1	< 0.1
Junior bonds			
Weighted average life (years)	5.00	5.00	5.00
Probability of default (bps)	< 0.1	< 0.1	< 0.1
Expected loss (bps)	< 0.1	< 0.1	< 0.1
Guarantee performance[f]			
Pr(draw on guarantee) (%)	3.2	0.1	0
Expected cost of guarantee ($K)	26.9	0.6	0
98^{th}-percentile of guarantee cost ($K)	407.7	0	0
99^{th}-percentile of guarantee cost ($mm)	1.0	0	0
Maximum cost of guarantee ($mm)	6.3	3.9	0

Abbreviations. PD: probability of default, IRR: internal rate of return, bps: basis points (1 bp = 0.01%), Pr: probability, K: thousands, mm: millions.

[a] For each scenario, 10 million Monte Carlo simulation paths are used.

[b] The expected loan default probability in this case is given by (C.2), and depicted by the dashed purple line in Figure 4.3.

[c] The expected loan default probability in this scenario is defined in (C.1), and represented by the solid orange line in Figure 4.3.

[d] The expected loan default probability for this scenario is defined in in (C.3), and represented by the dotted green line in Figure 4.3.

[e] For the definition of weighted average life, see (4.6).

[f] Cost of guarantee is derived using (4.7), where the guarantor's cost of capital is $R_g = 2\%$.

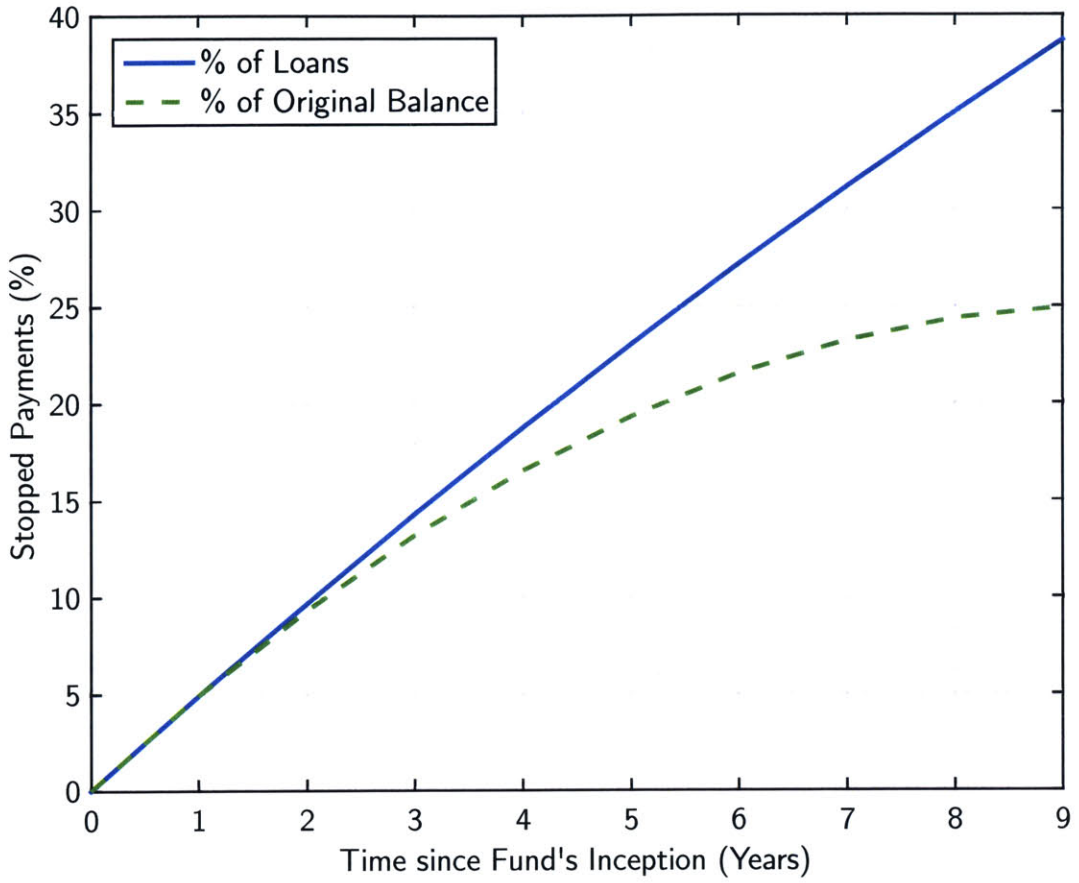

Figure 4.4. Cumulative losses of the portfolio over the fund's life both as a percentage of the number of initial HCLs and as a percentage of the original balance of HCLs in the portfolio.

performance (IRR = 15.8%, TCR = 161.6%), and the baseline case falls in between (IRR = 12.5%, TCR = 144.8%). The samples of the IRR distributions listed in Table 4.1 for these scenarios further support the same trend for their equity performances, as expected due to the trend in the probability of loan defaults across these scenarios (Figure 4.3a) and its impact on the equity performance. The probability density functions (pdfs) of the IRRs for the equity tranche of the HCL fund in each of the three scenarios are depicted in Figure 4.5. As is seen in Figure 4.5, the equity performance across all the considered scenarios is within an acceptable range for potential equity investors.

For the performance of HCL debt tranches, we evaluate the probability of default (PD) and the expected loss (EL) for the senior and junior tranches. Because of the guarantees on the bonds, both debt tranches demonstrate lowest possible PDs and ELs

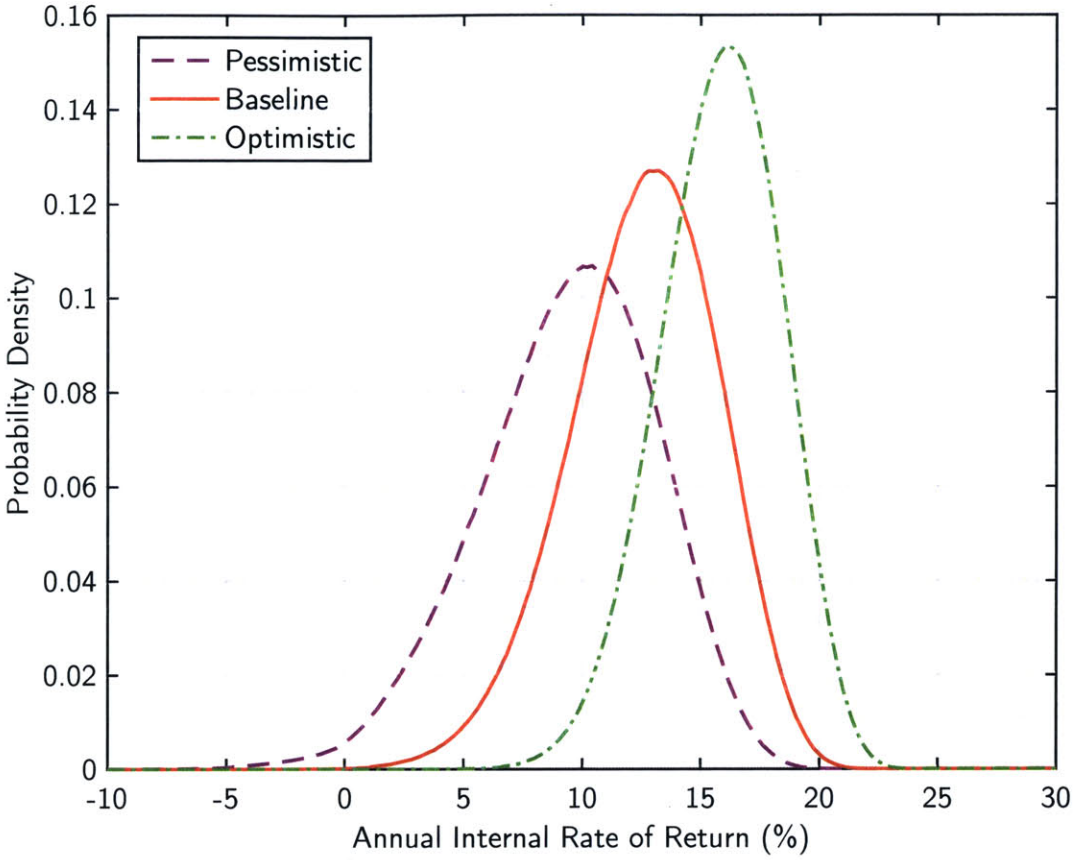

Figure 4.5. Probability density functions of the Internal rate of return (IRR) for the equity tranche of the HCL fund for the three scenarios of HCL defaults, namely, the pessimistic, baseline, and optimistic scenarios.

for all the three scenarios, making the bonds attractive to prospective investors. We also report the weighted average life (WAL) of the notes in Table 4.1, which is calculated as the following:

$$\text{WAL} = \left(\sum_{k=1}^{n} kP_k\right)\left(\sum_{k=1}^{n} P_k\right)^{-1}, \tag{4.6}$$

where P_k is the principal returned to the bondholders of each debt tranche in year k. Put simply, the weighted average life associated with each debt tranche denotes how long on average it takes for every invested dollar to be returned to the investors of that debt tranche. Because of credit enhancement techniques—especially, the cash flow diversion mechanisms—in place, the senior bonds have a shorter WAL than the scheduled 5-year WAL if the probability of loan default is large, i.e., in the baseline

and pessimistic scenarios in Table 4.1. For a detailed description of credit enhancement techniques, see Appendix C.1.

Lastly, we calculate the cost of guarantee (G), i.e., the cost incurred by the guarantor of the bonds, using the following equation:

$$G = \sum_{k=1}^{n} \frac{d_k}{(1+R_g)^k}, \tag{4.7}$$

where d_k is the amount drawn from the available guarantee in year k, and R_g is the assumed cost of capital for the guarantor ($R_g = 2\%$ for the numbers in Table 4.1, for a description of cost of capital, see Section 4.2).

Cost of guarantee, G in (4.7), is the present value of the drawn guarantee over the fund's life, and the expected cost of guarantee is simply its expected value. We report the probability that the guarantee is ever used to cover debt obligations during the entire life's fund, namely, $\Pr(G > 0)$, where $\Pr(\cdot)$ denotes probability.

In the pessimistic scenario, this probability is largest, standing only at 3.2%, while in the baseline and optimistic scenarios it is only 0.1% and 0, respectively. Not surprisingly, the pessimistic scenario has the highest expected cost of guarantee because of the higher loan default probabilities in this case. However, the cost of guarantee even in this worst-case scenario is only \$26.9K, a tiny fraction (0.6 bps) of the face value (amount) of the bonds, i.e., \$450mm, where bp stands for basis point and 1 bp = 0.01%. Lastly, the 98^{th} and 99^{th} percentiles of the cost of guarantee distribution as well as the maximum cost of guarantee, listed in Table 4.1, are clear testimonies to the negligible cost of guarantee compared to the face value of the bonds, i.e., \$450mm, across all the three scenarios.

■ 4.5 Aligning Interests: Moving Beyond "Pay for Performance"

We recommend that the insurance companies guarantee the bonds. On the one hand, by shifting a portion of the price to the patients, the insurance companies can, to some extent, reduce the large burden of upfront payments that they would face to cover all the warehoused patients in the current pricing environment. On the other hand, as discussed in the previous section, the cost of guarantee is a small fraction of the face value of the bonds; hence, in exchange for the savings the insurance companies make in this framework, they are highly motivated to guarantee the bonds so as to makes them attractive to the bond investors who are willing to lend at the lowest possible interest rates. This in turn reduces the financial burden of the loans on patients.

Another natural candidate for the guarantor of the bonds and/or investment in the junior bonds are pension funds whose financial liabilities move inversely with the mortality rates. Therefore, the guarantees or a position in the junior tranche would act as a natural hedge for their investment, reducing their exposure to changes in mortality rates.

The proposed structure can further be used in a pay-for-performance pricing environment, a recent focus of the medical profession and patient advocates. The performance of the equity tranche deteriorates if the death rates are relatively high (e.g., in the case where the therapy is not really curative). Pharmaceutical companies can, in theory, write some specialized contracts on the equity investment, not unlike stock put options, guaranteeing a bottom-line equity return for equity investors. With these contracts, the equity investors would be protected against high mortality rates and ineffective drugs, while the drug company would effectively disgorge a portion of their fees if these adverse events occur. This motivates pharmaceutical companies to produce more effective medications and to follow patients to ensure continued efficacy.

Lastly, because in our proposal, each patient is involved in the payments for their treatment, they would have some "skin in the game" (literally and figuratively), and will take better care of themselves, e.g., through more diligent adherence to their prescribed drug regimen, diet, exercise, and other actions that will "protect" their healthcare investment.

■ 4.6 Conclusion

In the current pricing environment for highly curative therapies and because of the lack of innovative payment options, the burden of the upfront payment, caused by the current high prices of such drugs and/or a large prevalence of the disease, makes it practically impossible for insurance companies to provide the largest number of patients with potentially life-saving therapies. To address this issue, we propose a new financing paradigm by having patients pay a portion of their drug's price over time through healthcare loans. We then use portfolio theory and securitization techniques to design a new fund to finance these loans, while keeping the cost of borrowing for the patients to the lowest possible levels. By estimating the post-treatment mortality rates of the patients and using a few statistical models for the default characteristics of these loans, we demonstrate a great potential for the proposed fund under practical conditions. Considering the extremely large burden of some prevalent diseases such as HCV that

already have curative therapies, taking action may not be a choice, but could soon become a necessity.

Chapter 5

Discussion and Conclusion

In this dissertation, we focus on three areas of pharmaceutical innovation and propose candidate solutions to improve the performance of each area. First, we investigate and improve upon a recently introduced financing vehicle for early-stage drug research and discovery. We use state-of-the-art financial engineering techniques to bring large sums of capital from global financial markets to early-stage pharmaceutical R&D projects. This method allows pharmaceutical companies to tap into a much larger pool of capital, and to engage in risk-sharing activities with a much larger pool of investors with a diverse set of risk appetites. We show that by slicing a drug R&D investment into appropriate sub-investments—each tailored to a different set of investors with distinct risk tolerances and associated return expectations—we can create a structure, where all the stakeholders benefit from the risk-sharing paradigm. This is feasible because we can reduce the inefficiency of current financing techniques used in the pharmaceutical industry by employing the proposed financing vehicle.

The second area on which we focus is the drug-regulatory environment, especially, drug-approval criteria for life-threatening illnesses. We propose a new quantitative framework which can objectively incorporate the severity of disease into the drug-approval criteria and into the required balance between benefits and risks in order for a drug to be approved. We use Bayesian decision analysis to design a model to explicitly take into account the severity of disease at the drug-approval stage and we then use the Global Burden of Disease Study 2010 to set the parameter values in the model [43]. Using the proposed framework, we demonstrate that the current drug-regulatory criteria may be too conservative for the most life-threatening diseases while too permissive for milder conditions.

In the last part of this dissertation, we look at the burdensome aggregate cost, that would be incurred by insurance companies were they to cover the largest patient population for some of the recently approved curative therapies. To support the reward

system for pharmaceutical companies and to keep them motivated to discover truly innovative medicine, we propose a new financing paradigm for these highly curative therapies. This new financing method is tailored to the current environment as the industry moves closer to personalized medicine and more innovative solutions are needed to cover the costs of those therapies.

Appendix A

Appendix to Chapter 2

In focusing our study on the implications of dynamic leverage in Chapter 2, we have not addressed a number of practical issues regarding the launching and management of a megafund. Some of these issues have been addressed in other articles [13, 14, 20–25], but in the interest of completeness and convenience, we summarize the most relevant of these issues here in the form of "frequently asked questions".

Q: What is the intuition for the advantages of dynamic leverage, should you also consider dynamic equity (new equity issues), and is it related to the tax deductibility of interest payments?

A: The intuition for the advantages of dynamic leverage is that the cost of equity for biotech businesses is in the 15% to 30% range whereas the cost of debt is assumed to be in the 5% to 8% range, and this lower cost of capital allows equity-holders to benefit from the spread. Moreover, we adjust the level of debt every period to make sure that bondholders are not being over-compensated when the default risk is too little. If dynamic equity could enhance the model, we could incorporate that as well. However, based on our analysis, as the portfolio matures, it is effectively over-equitized (we return equity to equity investors, rather than needing additional equity) and to achieve an optimal capital structure, all additional capital should be debt. Tax deductibility would be an additional benefit.

Q: Does dynamic leverage not increase the risk to equity holders?

A: There is no doubt in that the equity risk is higher with debt than without debt (the risk of the megafund's equity tranche is higher than that of the equity-only portfolio). However, in Chapter 2, we assert that the equity return of the megafund with dynamic leverage is better than the equity return of the megafund without dynamic leverage. This is true because, as mentioned above, we raise debt to pay distributions to equity, hence, the return on equity is smoother than the case where there is no debt adjustment noting that the equity is the most junior tranche. We should note that we are able to

achieve this performance enhancement with no added risk for equity-holders, compared to the megafund without dynamic leverage, because the bondholders are not over-compensated when the portfolio performs well.

Q: Would existing bondholders not object to new debt issues, and would it not be a challenge to get their permission each time leverage is dynamically adjusted?

A: First, because we only increase the amount of junior debt—which is less senior than the senior debt (i.e., has a lower priority for default payments, etc.)—senior bondholders should not object to this mechanism. However, the point made about the junior bondholders losing a part of their investment value because of maintaining a constant default probability, is exactly why we propose using the dynamic leverage method in the first place. Otherwise, in situations where the portfolio performs well, the bondholders will be paid yields that are much higher than what the market suggests for those new (and lower) default rates. As we stated above, the level of debt is adjusted every period to make sure that bondholders are not being over-compensated when the default risk is too little.

In addition, the dynamic leverage structure would be transparent to the initial lenders/investors and they would participate with the knowledge that additional debt can be added under specific conditions. The loan documents or securitization indenture would incorporate a contractual ability to increase leverage if certain conditions are met. The assumption is that those conditions are described by the model developed. An analog in the securitization world is a master trust for credit card debt. As more receivables are added to the pool, additional debt may be issued. This debt may be structured to have identical terms to the existing debt, or may be "tranched" into different rating classes that have terms equal to previously issued bonds. Another example is a bank loan facility secured by a pool of assets. As the pool of assets increases, additional funds may be drawn from the lenders. In more complex structures, this may entail several tranches of debt of different seniority. In each case, the ability to issue or draw more debt depends on the specific conditions of the loan agreements or indentures.

Q: How would a megafund be managed?

A: A megafund is not merely a passive investment vehicle, but must be professionally managed by a team of experts in life-sciences investing, i.e., former biotech venture capitalists or financial analysts of the bio-pharmaceutical industry. An existing example of such a structure is a drug royalty investment company such as Royalty Pharma. The

example is not exact because Royalty Pharma currently does not invest in early-stage assets.

Q: Is it not difficult to value assets?

A: Drug royalty investment companies often hire scientific consultants as well as third-party valuation groups to provide valuations. Similar approaches are used in securitization strategies in other industries such as credit-card receivables, student and auto loans, film royalties, music royalties, etc. Among securitizations of future cash flows, film royalty securitizations are the closest to the drug megafund structure. In those deals, the valuation of the individual films is not known in advance and the securitization is based on the probability that the cashflows will meet or exceed historical results.

Q: What protects the megafund from adverse selection, i.e., being offered the poorest assets by drug developers?

A: Misaligned incentives can indeed result in the adverse selection of projects, in which a megafund is shown the least attractive assets and never has the opportunity to acquire the most attractive ones. This potential bias is the reason that a megafund has to be managed by a team of experienced life-sciences investment professionals. However, this bias can be addressed if the drug developer or asset owner has sufficient "skin in the game," i.e., a continuing ownership interest in the compound even after the megafund acquires it. In this way, the interests of both drug developer/asset owner and investor will be aligned. Currently, the parameters for success rates that we have used in our simulations are estimated using databases that include compounds licensed to different companies at some point in their development cycle. Hence, if partnership strategies have any negative effect on the success rate of compounds, this effect is likely to be reflected in the data used to estimate our model's parameters. A megafund would certainly have detailed incentives, operating plans and governance controls, all outside the scope of our paper, to ensure that the risk of adverse selection is minimized and contained.

Q: The simulations employ parameters that are based on historical averages; are these likely to apply to current and future projects?

A: There are two issues implicit in this question: (a) do historical averages reflect potential incentive issues that would reduce success rates?; and (b) are historical averages representative of future performance? We have addressed (a) in our previous answer, and believe that historical figures already incorporate the impact of adverse selection, and with sufficiently skilled investment management, a megafund should be able to

achieve such rates or better. However, (b) is a broader concern, especially given the scientific progress that has occurred in just the past five years in biomedicine. It is quite likely that historical success rates do not reflect the recent breakthroughs in bioinformatics, companion diagnostics, and the various "omics." These innovations are likely to make our historical success rates conservative estimates of what a megafund could realize today. Moreover, even though history is an imperfect guide for the future, investors have no choice but to use historical data as the basis for their investment decisions, and clearly they do so on a regular basis, as the popularity of various commercial drug development databases from EvaluatePharma, Informa, Thomson Reuters, and other vendors suggest.

Q: Are your parameter assumptions realistic?

A: With respect to our biomedical parameter assumptions such as clinical success rates, we have no domain-specific expertise and therefore rely on industry publications to select these parameters. However, open-source software used to perform these simulations is available online and readers are encouraged to re-run our simulations using their own parameter values to check the robustness of our results (see `http://alo.mit.edu/wp-content/uploads/2015/06/RBOtoolbox_final.zip` and [13] for further details).

The simulations conducted in this study focus specifically on orphan diseases, and the parameters we use are specifically chosen for these types of assets [14]. A recent study using a live portfolio of orphan drug projects managed by the National Center for Advancing Translational Sciences (NCATS) suggests that our parameters are conservative [21].

With respect to our financial parameters, the assumed yields for our debt tranches are actually quite a bit higher than current market rates, given their default rates and expected losses. For example, as of Friday November 13, 2015, the composite corporate bond yield for a 20-year AAA corporate bond is 3.92%[1] and the cumulative average default probability for AAA bonds 10 years after issuance is 0.77% or 77bps.[2] Our simulations assume that such a bond yields 5.00% and the simulated default rate in our analysis is less than 0.1bps hence our assumptions are considerably more conservative than current market values. In addition, a cushion was added to incorporate additional costs such as hedging, fees and expenses of the financing, commitment fees, etc.

[1] See `http://finance.yahoo.com/bonds/composite_bond_rates`.

[2] See Table 24 of `https://www.nact.org/resources/2014_SP_Global_Corporate_Default_Study.pdf`.

Appendix B

Appendix to Chapter 3

In this Appendix, we derive the expected-cost-minimizing critical value and sample size in Section B.1, and in Section B.2 we show how to impute the costs of Type I and Type II errors implicit in any one-sided fixed-sample test of a given size and power under the assumption that it is BDA-optimal.

■ B.1 Expected Cost Optimization

We determine the optimal sample size and the critical value for the fixed-sample test, $\texttt{fxd}(n, \lambda_n)$, by minimizing its expected cost in (3.5) over all possible values for n and λ_n. Keeping the sample size n fixed, the critical value λ_n that minimizes the expected cost, $C(\texttt{fxd}(n, \lambda_n))$ in (3.5), can be determined by setting the partial derivative of the expected cost, with respect to λ_n, to zero:

$$\frac{\partial}{\partial \lambda_n} C(\texttt{fxd}(n, \lambda_n))\Big|_{\lambda_n = \lambda_n^*} = N p_0 c_1 \left[-\phi\left(-\lambda_n^*\right) + \bar{c}_2 \phi\left(\lambda_n^* - \delta_0 \sqrt{\mathcal{I}_n}\right)\right] = 0, \qquad \text{(B.1)}$$

where ϕ is the probability density function (pdf) of a standard normal random variable, i.e., $\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}x^2\right)$. Now, solving (B.1) for λ_n^* yields:

$$\lambda_n^* = -\frac{1}{\delta_0 \sqrt{\mathcal{I}_n}} \log\left(\bar{c}_2\right) + \frac{\delta_0 \sqrt{\mathcal{I}_n}}{2}, \qquad \text{(B.2)}$$

where log is the natural logarithm and $\mathcal{I}_n = \frac{n}{2\sigma^2}$. By calculating the second derivative of the expected cost in (3.5), it is straightforward to prove that λ_n^*, given by (B.2), indeed minimizes the expected cost, $C(\texttt{fxd}(n, \lambda_n))$. Assuming $p_0 = p_1 = 0.5$, if Type I and Type II costs were equal, it is clear that the optimal critical value should be the midpoint of the means of the Z-statistic under the two hypotheses, hence the existence of the term $\frac{\delta_0 \sqrt{\mathcal{I}_n}}{2}$ in (B.2). However, in a general case, where the two costs are distinct,

the first term in (B.2) plays the role of a correction term, and adjusts the optimal critical value to incorporate the difference between Type I and Type II costs.

Given a specific value of $\bar{c}_2$, the optimal critical value, i.e., λ_n^* in (B.2), can be considered a function of the sample size. The behavior of this function over different sample sizes for three values $c_2 = 0.01, 0.07, 0.34$, corresponding to $\bar{c}_2 = 0.2, 1, 5$, respectively, is depicted in Figure B.1, where the alternative hypothesis corresponds to $\delta_0 = \frac{\sigma}{8}$. The conventional critical value, regularly used for one-sided tests, i.e., $z_\alpha = \Phi^{-1}(1-\alpha) = 1.96$ for $\alpha = 2.5\%$, is also drawn in Figure B.1 for comparison. It is observed that, in all of these cases, the optimal critical value changes with the sample size contrary to the classical critical value, which is independent of the sample size.

Now, if we assume equally likely hypotheses, i.e., $p_0 = p_1 = 0.5$, the parameter $\bar{c}_2$ becomes the ratio of Type II cost to Type I cost, which must be larger for life-threatening diseases than for mild diseases, as discussed in Section 3.4. In other words, the parameter $\bar{c}_2$ can be considered as a normalized indicator of the severity of the target disease. The more dangerous the disease, the higher the value of $\bar{c}_2$ should be, and the larger chance we should give to an effective drug to be approved. Therefore, across all sample sizes in Figure B.1, by increasing the value of $\bar{c}_2$, the optimal critical value becomes smaller and moves toward the mean of the Z-statistic under the null hypothesis, namely, the constant zero line. In other words, the optimal critical value becomes less conservative as the importance of Type II cost relative to Type I cost increases, modeling a more life-threatening disease. This explains why the red line lies completely below the green line and the green curve is below the blue line. If $\bar{c}_2$ is large enough, the optimal critical value may even cross the zero line and become negative, e.g., if $\bar{c}_2 = 5$, λ_n^* becomes negative over sample sizes smaller than 779. Finally, for $\bar{c}_2 < 1$, implying a larger weight for the Type I cost, corresponding to mild diseases, the behavior of the optimal critical value is qualitatively different from the other two cases, in which the optimal critical value is monotonically increasing in the sample size.

Using the optimal critical value in (B.2), the size, α, and the power of the test at the alternative hypothesis, $1-\beta$, are given by

$$\alpha = \Phi\left(\frac{1}{\delta_0\sqrt{\mathcal{I}_n}}\log(\bar{c}_2) - \frac{\delta_0\sqrt{\mathcal{I}_n}}{2}\right), \tag{B.3}$$

$$1-\beta = \Phi\left(\frac{1}{\delta_0\sqrt{\mathcal{I}_n}}\log(\bar{c}_2) + \frac{\delta_0\sqrt{\mathcal{I}_n}}{2}\right). \tag{B.4}$$

Next, for a given n, the expected cost obtained by using the optimal critical value, λ_n^*

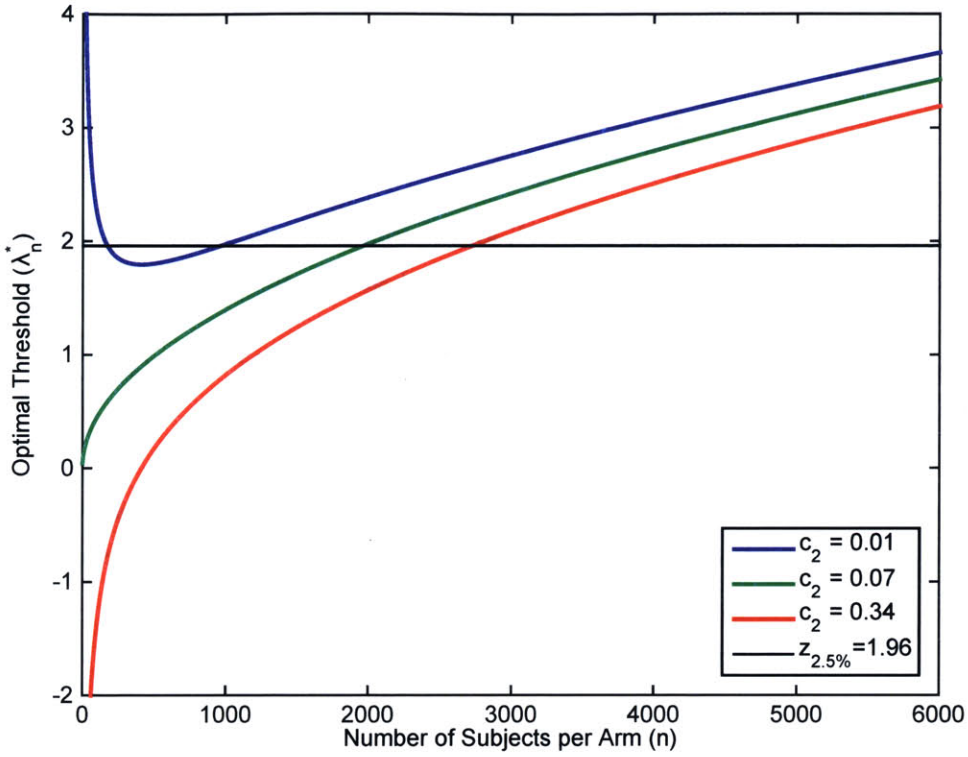

Figure B.1. The optimal critical value as a function of the number of subjects per arm for three different diseases.

Note. The severity of disease is denoted as c_2 where $c_2 = 0.01$, corresponding to $\bar{c}_2 = 0.2$, represents mild severity. Medium severity corresponds to $c_2 = 0.07$, or equivalently $\bar{c}_2 = 1$, and life-threatening disease is denoted by $c_2 = 0.34$, corresponding to $\bar{c}_2 = 5$. The constant line with the height $z_{2.5\%} = 1.96$ (thin black line) is also drawn for comparison.

in (B.2), can be calculated by substituting (B.2) into (3.5) and is given by

$$C(\texttt{fxd}(n, \lambda_n^*)) = p_0 c_1 \left[N\Phi(-\lambda_n^*) + N\bar{c}_2\Phi(\lambda_n^* - \delta_0\sqrt{\mathcal{I}_n}) + n(1 + \gamma N\bar{c}_2) \right], \qquad \text{(B.5)}$$

where the optimal sample size should be determined to minimize this expected cost over all possible sample sizes. Let us consider a continuum of values, rather than discrete values, for the sample size, n, and take the partial derivative of the expected cost in

(B.5) with respect to n as the following:

$$\frac{\partial}{\partial n} C(\texttt{fxd}(n, \lambda_n^*)) = \left(\frac{\partial}{\partial n}\lambda_n^*\right) N\left[-\phi(-\lambda_n^*) + \bar{c}_2\phi(\lambda_n^* - \delta_0\sqrt{\mathcal{I}_n})\right] - \left(\delta_0\frac{\partial}{\partial n}\sqrt{\mathcal{I}_n}\right) N\bar{c}_2\phi(\lambda_n^* - \delta_0\sqrt{\mathcal{I}_n}) + (1 + \gamma N\bar{c}_2), \tag{B.6}$$

where the first line is proportional to (B.1), and therefore, equal to zero. By simplifying (B.6) and setting it to zero to evaluate the optimal sample size n^*, we have:

$$\frac{\partial}{\partial n} C(\texttt{fxd}(n, \lambda_n^*))\Big|_{n=n^*} = \left[(1 + \gamma N\bar{c}_2) - \left(\delta_0\frac{\partial}{\partial n}\sqrt{\mathcal{I}_n}\right) N\bar{c}_2\phi(\lambda_n^* - \delta_0\sqrt{\mathcal{I}_n})\right]\Big|_{n=n^*} = 0. \tag{B.7}$$

Now, if we define $x^* \triangleq \frac{1}{2}\left(\delta_0\sqrt{\mathcal{I}_{n^*}}\right)^2$, and rearrange terms in (B.7), then x^* can be represented as the fixed point of a function, g, i.e., $x^* = g(x^*)$. This function is given by

$$g(x) = A\exp\left(-\frac{1}{2}\left(\frac{\log^2(\bar{c}_2)}{x} + x\right)\right), \tag{B.8}$$

where

$$A = \frac{N^2}{16\pi}\left(\frac{\delta_0^2}{2\sigma^2}\right)^2\left[\frac{\bar{c}_2}{(1 + \gamma N\bar{c}_2)^2}\right] \tag{B.9}$$

and $\bar{c}_2 = \frac{p_1 c_2}{p_0 c_1}$ as defined earlier. Now, if N is large enough to make $\gamma N\bar{c}_2$ much larger than 1, A becomes independent of N. Since the exponential function in g is independent of N as well, we observe the insensitivity of the optimal sample size to the prevalence of disease, N, in the case of large burden of disease, i.e., large $C_2 = Nc_2$. In the following, we revisit the three cases, for which the optimal critical value is drawn in Figure B.1. Let us consider, for all these cases, a target population of $N = 500{,}000$ patients, an alternative hypothesis associated with $\delta_0 = \frac{\sigma}{8}$ and equal prior probabilities for the two hypotheses, i.e., $p_0 = p_1$. We consider $\bar{c}_2 = 0.2$ (equivalently $c_2 = 0.01$) corresponding to an innocuous disease, $\bar{c}_2 = 1$ (equivalently $c_2 = 0.07$) representing a disease with medium severity, and $\bar{c}_2 = 5$ (equivalently $c_2 = 0.34$) corresponding to a life-threatening disease. By using (B.8), we first determine the optimal sample size, then substitute this n^* into (B.2) to determine the optimal critical value, and finally, using (B.3) and (B.4), we calculate the size of the optimal tests, and their power for the alternative hypothesis. The results are tabulated in Table B.1.

Table B.1. The optimal sample size, critical value, size, and statistical power for three trials, each designed to test a treatment targeting a disease with a different severity.

Severity	Optimal Sample Size	Optimal Critical Value	Size (%)	Power (%)
0.01	2,719	2.654	0.40	97.47
0.07	2,236	2.090	1.83	98.17
0.34	1,534	1.266	10.28	98.59

For the three listed trials, the size of the target population is $N = 500{,}000$ and the alternative hypothesis corresponds to $\delta_0 = \frac{\sigma}{8}$, for which the power is reported.

In the following section, we employ the cost model proposed in Section 3.4 and the results of this section to determine the implicit costs in the current standards of clinical trials.

■ B.2 Imputing the Cost of Type I and Type II Errors

We consider a typical one-sided fixed-sample test and assume that it is a BDA-optimal test in our framework using some unknown normalized cost parameter $\bar{c}_2$ and unknown prevalence, N, and infer these parameters for the trial. The FDA regulations require that one-sided tests have at most 2.5% probability of Type I error. For this current standard, it is easy to see that the critical value in a fixed-sample test, on the Z-scale, is

$$\lambda_n^* = z_\alpha \triangleq \Phi^{-1}(1-\alpha), \tag{B.10}$$

where $\alpha = 2.5\%$ and hence, $z_\alpha = 1.960$. Also, because the Type II error associated with δ_0 is equal to β, we have:

$$\beta = \Phi(\lambda_n^* - \delta_0\sqrt{\mathcal{I}_n}) \quad \Rightarrow \quad z_\beta \triangleq \Phi^{-1}(1-\beta) = \delta_0\sqrt{\mathcal{I}_n} - \lambda_n^* = \delta_0\sqrt{\mathcal{I}_n} - z_\alpha. \tag{B.11}$$

Substituting (B.10) and (B.11) into (B.2) gives us:

$$z_\alpha = (z_\alpha + z_\beta)^{-1}\log\left(\frac{p_0 c_1}{p_1 c_2}\right) + \frac{z_\alpha + z_\beta}{2} \quad \Rightarrow \quad \log\left(\frac{p_0 c_1}{p_1 c_2}\right) = \frac{z_\alpha^2 - z_\beta^2}{2}. \tag{B.12}$$

This yields the ratio of Type II cost to Type I cost as:

$$\bar{c}_2 = \exp\left(\frac{z_\beta^2 - z_\alpha^2}{2}\right) = \exp\left(\frac{1}{2}\left(\frac{\delta_0^2}{2\sigma^2}\right)n - z_\alpha\sqrt{\frac{\delta_0^2}{2\sigma^2}n}\right). \tag{B.13}$$

Note that the cost ratio depends on the number of subjects recruited in the trial. However, in our model, this cost ratio is an exogenous variable which is related to the severity of the targeted disease, the state of current therapies for the disease, and the side effects of the drug. Therefore, the ratio should not depend on the sample size.

Now, in classical hypothesis testing, $\lambda_n^* = z_\alpha$, which is independent of the sample size. Using this fact, we can further simplify the conditions for optimal sample size by noting that the optimal sample size n^*, is the integer value n, for which:

$$\frac{C(\texttt{fxd}(n+1, \lambda_{n+1}^*))}{C(\texttt{fxd}(n, \lambda_n^*))} \geq 1 \quad \text{and} \quad \frac{C(\texttt{fxd}(n-1, \lambda_{n-1}^*))}{C(\texttt{fxd}(n, \lambda_n^*))} > 1. \tag{B.14}$$

By expanding the left-hand side of (B.14), we have:

$$\begin{aligned}
\frac{C(\texttt{fxd}(n+1, \lambda_{n+1}^*))}{C(\texttt{fxd}(n, \lambda_n^*))} &= \frac{N\Phi(-z_\alpha) + N\bar{c}_2\Phi\Big(z_\alpha - (z_\alpha + z_\beta)\sqrt{\frac{n+1}{n}}\Big) + (n+1)(1+\gamma N\bar{c}_2)}{N\Phi(-z_\alpha) + N\bar{c}_2\Phi(-z_\beta) + n(1+\gamma N\bar{c}_2)} \\
&= 1 + \frac{(1+\gamma N\bar{c}_2) - N\bar{c}_2\left[\Phi(-z_\beta) - \Phi\Big(-(z_\alpha+z_\beta)\left[\sqrt{1+\frac{1}{n}} - 1\right] - z_\beta\Big)\right]}{C(\texttt{fxd}(n, \lambda_n^*))} \\
&= 1 + \frac{(1+\gamma N\bar{c}_2) - N\bar{c}_2 \Pr(-z_\beta - \epsilon_1 < Z \leq -z_\beta)}{C(\texttt{fxd}(n, \lambda_n^*))} \geq 1,
\end{aligned} \tag{B.15}$$

where $Z \sim \mathcal{N}(0,1)$ and $\epsilon_1 = (z_\alpha + z_\beta)\left[\sqrt{1+\frac{1}{n}} - 1\right]$. Furthermore, by expanding the second inequality in (B.14), we have:

$$\begin{aligned}
\frac{C(\texttt{fxd}(n-1, \lambda_{n-1}^*))}{C(\texttt{fxd}(n, \lambda_n^*))} &= \frac{N\Phi(-z_\alpha) + N\bar{c}_2\Phi\Big(z_\alpha - (z_\alpha + z_\beta)\sqrt{\frac{n-1}{n}}\Big) + (n-1)(1+\gamma N\bar{c}_2)}{N\Phi(-z_\alpha) + N\bar{c}_2\Phi(-z_\beta) + n(1+\gamma N\bar{c}_2)} \\
&= 1 + \frac{-(1+\gamma N\bar{c}_2) + N\bar{c}_2\left[\Phi\Big((z_\alpha+z_\beta)\left[1-\sqrt{1-\frac{1}{n}}\right] - z_\beta\Big) - \Phi(-z_\beta)\right]}{C(\texttt{fxd}(n, \lambda_n^*))} \\
&= 1 + \frac{-(1+\gamma N\bar{c}_2) + N\bar{c}_2 \Pr(-z_\beta < Z \leq -z_\beta + \epsilon_2)}{C(\texttt{fxd}(n, \lambda_n^*))} > 1,
\end{aligned} \tag{B.16}$$

where $\epsilon_2 = (z_\alpha + z_\beta)\left[1 - \sqrt{1-\frac{1}{n}}\right]$. Next, combining (B.15) and (B.16) yields:

$$\Pr(-z_\beta - \epsilon_1 < Z \leq -z_\beta) \leq \frac{1+\gamma N\bar{c}_2}{N\bar{c}_2} < \Pr(-z_\beta < Z \leq -z_\beta + \epsilon_2). \tag{B.17}$$

Now, both ϵ_1 and ϵ_2 can be well-approximated by $\epsilon_1 \approx \epsilon_2 \approx \frac{z_\alpha + z_\beta}{2n} = \left(\frac{1}{2}\sqrt{\frac{\delta_0^2}{2\sigma^2}}\right)\sqrt{n^{-1}}$. Hence, for a relatively large sample size,[1] n, we can simplify (B.17) to yield:

$$\frac{1+\gamma N\bar{c}_2}{N\bar{c}_2} \approx \left(\frac{z_\alpha + z_\beta}{2n}\right)\left[\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}z_\beta^2\right)\right] = \frac{1}{2\sqrt{2\pi n}}\sqrt{\frac{\delta_0^2}{2\sigma^2}}\exp\left(-\frac{1}{2}z_\beta^2\right). \quad \text{(B.18)}$$

Now, we multiply the result in (B.18) by (B.13) to get the following ratio:

$$\frac{1+\gamma N\bar{c}_2}{N} \approx \frac{1}{2\sqrt{2\pi n}}\sqrt{\frac{\delta_0^2}{2\sigma^2}}\exp\left(-\frac{1}{2}z_\alpha^2\right), \quad \text{for large } n. \quad \text{(B.19)}$$

Therefore, for a balanced two-arm fixed-sample test with n subjects per arm and a size of α, i.e., $\texttt{fxd}(n, z_\alpha)$, which has a power of $1-\beta$ at δ_0, we can estimate the normalized Type II cost and prevalence of the disease as:

$$\bar{c}_2 = \exp\left(\frac{z_\beta^2}{2} - \frac{z_\alpha^2}{2}\right) = \exp\left(\frac{1}{2}\left(\frac{\delta_0^2}{2\sigma^2}\right)n - z_\alpha\sqrt{\frac{\delta_0^2}{2\sigma^2}n}\right), \quad \text{(B.20)}$$

$$N \approx \frac{1}{\frac{1}{2\sqrt{2\pi n}}\sqrt{\frac{\delta_0^2}{2\sigma^2}}\exp(-\frac{1}{2}z_\alpha^2) - \gamma\exp(\frac{z_\beta^2 - z_\alpha^2}{2})}, \quad \text{(B.21)}$$

where N is the size of the target population of the drug under test.

To put the results in (B.20) and (B.21) into perspective, let us assume that a fixed-sample test is required with size $\alpha = 2.5\%$, and a power of $1-\beta = 85\%$ for an alternative hypothesis corresponding to $\delta_0 = \frac{\sigma}{8}$. Using the classical hypothesis-testing calculations, this leads to a sample size of $n = 1{,}150$ which meets the FDA's criterion. Now let us assume a non-informative prior, i.e., $p_0 = p_1 = 0.5$. For this trial, we get the following severity and prevalence:

$$c_2 = 0.02, \quad N = 15{,}119. \quad \text{(B.22)}$$

Having obtained a severity equal to 0.02 in (B.22), we can conclude that the current standards for clinical trials are optimal for testing only innocuous diseases, as discussed in Sections 3.2 and 3.4, and cannot be optimal for more life-threatening diseases like pancreatic cancer. In general, as observed in Table B.1, for a given target population, the test should become less conservative (its critical value should become smaller) and

[1] For the numerical example given at the end of this section, if the number of recruited patients is more than 100, n is large enough for this approximation to hold.

Table B.2. The required sample size, implied severity, and prevalence of the target disease for four conventional trials.

Power (%)	Required Sample Size	Implied Severity	Implied Prevalence ($\times 1000$)
80	1,005	0.01	13.68
85	1,150	0.02	15.12
90	1,345	0.02	17.51
95	1,664	0.04	24.60

Each trial corresponds to a different power for the alternative hypothesis, namely, $1 - \beta = 80\%, 85\%, 90\%, 95\%$, and all the trials have a size of $\alpha = 2.5\%$. For all the trials, the alternative hypothesis corresponds to $\delta_0 = \frac{\sigma}{8}$.

the sample size should shrink as the severity of the disease increases to avoid exposure to inferior treatment during the trial. Now, maintaining all our assumptions and only changing the power of the test for the alternative hypothesis, we get different values for the required sample size, implied severity, c_2, and prevalence, N. We have reported the results for four different power levels for the alternative hypothesis, namely, $1 - \beta = 80\%, 85\%, 90\%$, and 95%, in Table B.2. As observed in Table B.2, all the implied severity values for these classical tests are too small, especially, for a high power level of 95% where the implied severity is only 0.04 (last row in Table B.2). These small numbers underscore the fact that the current standards of clinical trials are quite conservative and not suitable for terminal illnesses with no effective treatment.

Appendix C

Appendix to Chapter 4

C.1 Credit Enhancement Techniques

In this section, we describe two of the most commonly used credit enhancement techniques employed in securitization vehicles.

C.1.1 Interest Coverage Test

After making all the bond payments in each year, we calculate the ratio of the expected cash, to be received in the next year from the loan repayments, to the payments scheduled for the senior bonds in the next year. If this ratio is lower than 125%, we divert cash flows to the senior bonds, and reduce their outstanding principal such that this interest coverage ratio becomes lower than 125%, available cash permitting. After having the interest coverage ratio for the senior tranche pass, we calculate a similar ratio with the same numerator but a different denominator for the junior bonds. The denominator this time is the sum of the bond payments scheduled for both senior and junior notes. If this new ratio is less than 110%, we start paying the principal of the senior bonds down until the ratio is back in line. If all the principal of the senior bonds is paid, and the ratio is still higher than 110%, we pay down a portion of the junior bonds' principal to bring back the ratio to below 110%. After having both of these interest coverage tests pass, we move on to the overcollateralization test explained next.

C.1.2 Overcollateralization Test

For this test, instead of comparing the expected cash flow to be received by the portfolio with the scheduled payments to be made by the portfolio, we compare the outstanding amount of the assets in the portfolio, i.e., the loans, with the outstanding principals of the junior and senior tranches. Similar to the interest coverage ratio, for each tranche, the denominator is the sum of the outstanding principals of that tranche and all the

tranches senior to it. The numerator is simply the outstanding amount of the assets that the portfolio currently holds, i.e., the loans that have not stopped their payments. We use the same thresholds for cash diversion to the two tranches.

■ C.2 Loan Default Models

In this section, we propose a simple statistical model for two borrowers' default probabilities as drawn in Figure C.1. Suppose Alice (A) and Bob (B) are two borrowers, with two debt-payment-to-income ratios (DPI), defined as the ratio of the annual loan payments ($6,700 in our simulations) to the borrower's annual income. These DPIs derive the leftmost blocks in Figure C.1, where each borrower's initial expected default probability is uniquely determined by the input to the block, namely, his/her DPI. In each year, a common stochastic factor, modeling all the movements in economic conditions and wages, modulates both the initial expected default probabilities to yield the borrowers' final expected probabilities of default that are correlated with each other. We assume that the magnitude of the modulating signal is uniformly distributed over the interval $[75\%, 125\%]$, and that the magnitude in one year is independent of the magnitudes in other years. Finally, these correlated expected default probabilities feed into two separate beta random variable generators (the rightmost blocks in Figure C.1), creating two correlated beta random variables, each of which has a mean equal to its corresponding expected default probability (the input to the block).

Our goal in this section is to estimate the dependence of each initial expected default probability on its corresponding debt-payment-to-income ratio in the leftmost blocks in Figure C.1. We propose three different parametric models for this dependence, each of which is characterized by two parameters, and we then use student loan data to calibrate the parameters of each model. The baseline model, a generalization of the commonly used logistic model, is given by:

$$E[\text{PD}] = \begin{cases} \frac{1}{1+\exp\left(-\beta\Phi^{-1}(\text{DPI})+\alpha\right)}, & 0 \leq \text{DPI} \leq 1 \\ \\ 1, & 1 < \text{DPI} \end{cases}, \tag{C.1}$$

where $E[\text{PD}]$ is the expected loan default probability, $\Phi^{-1}(\cdot)$ is the inverse of the cumulative distribution function (CDF) of the standard normal random variable, and $\alpha \geq 0$ and $\beta \geq 0$ are the model parameters.

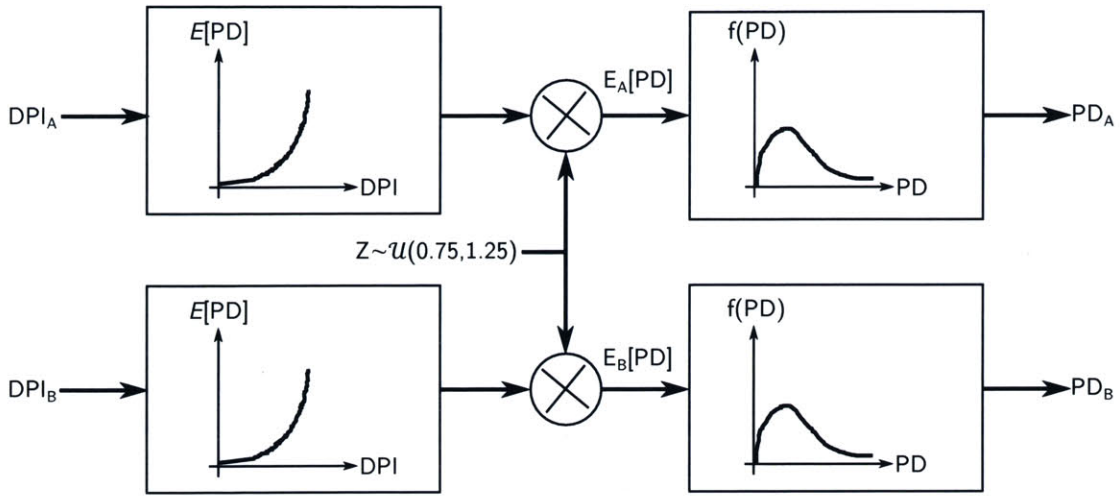

Figure C.1. The statistical model used to generate loan default probabilities.

Note. The object $\mathcal{U}(a, b)$ denotes the uniform probability distribution over the interval (a, b). DPI and PD stand for debt-payment-to-income ratio and probability of default, respectively.

The second model, referred to as the pessimistic model, is characterized by two parameters $\alpha \geq 0$ and $\beta > 0$ as the following:

$$E[\text{PD}] = \begin{cases} \exp\left(-\left(\frac{\text{DPI}^{-1}-1}{\beta}\right)^{\alpha}\right), & 0 \leq \text{DPI} \leq 1 \\ \\ 1, & 1 < \text{DPI} \end{cases}. \tag{C.2}$$

Lastly, the third model, labeled as the optimistic model, has two parameters $\alpha \geq 0$ and $\beta > 0$, and is given by:

$$E[\text{PD}] = \begin{cases} \exp\left(-\left(\frac{1-\text{DPI}}{\beta}\right)^{\alpha}\right), & 0 \leq \text{DPI} \leq 1 \\ \\ 1, & 1 < \text{DPI} \end{cases}. \tag{C.3}$$

The pessimistic and optimistic models are two variations of the Weibull distribution, for which the independent variable is $\text{DPI}^{-1} - 1$ and $1 - \text{DPI}$, respectively.

We use the 3-year cohort default rates of the federal student loans for the fiscal years 2009–2011 made available by the U.S. Department of Education[1] to estimate

[1]For a detailed description of the cohort default rates and their calculation, see .

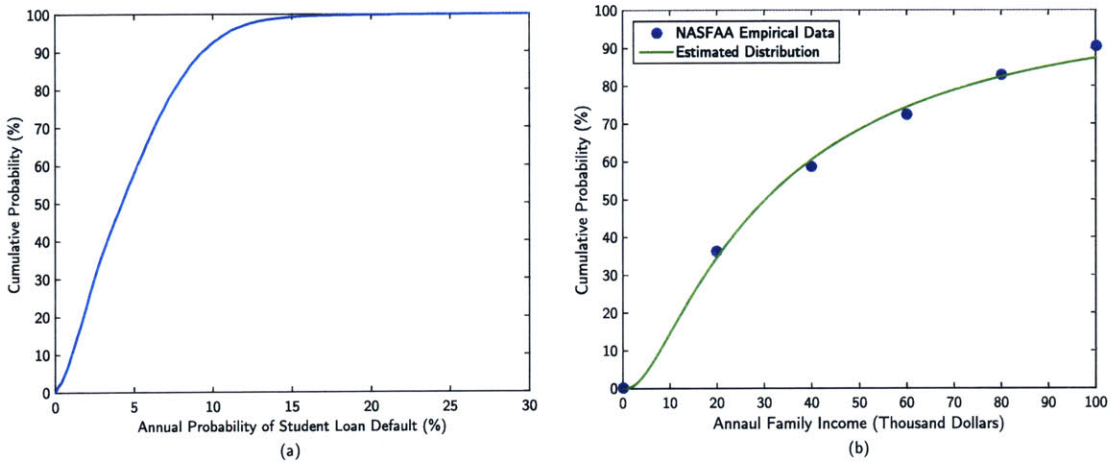

Figure C.2. The cumulative distribution of the annual probability of default for federal student loans, and (b) the estimated cumulative distribution of family-income for student loan borrowers along with the empirical numbers reported by the National Association of Student Financial Aid Administrators.

a probability distribution for the annual default probability on federal student loans. Since the reported default rates correspond to a 3-year time window and our goal is to estimate the annual default probabilities, we annualize the default rates using $p_1 = 1-(1-p_3)^{\frac{1}{3}}$, where p_1 and p_3 denote the annual and 3-year default rates, respectively. The estimated distribution of the annual default probabilities is drawn in Figure C.2a.

To estimate the distribution of family-income for the student loan borrowers, we then use the family-income levels of the students who received subsidized and unsubsidized Stafford loans[2] in the year 2007–08, reported by the National Association of Student Financial Aid Administrators (NASFAA) [83]. We fit a lognormal distribution to these empirical data to get a continuous cumulative distribution function (CDF) for the family-income distribution of the student loan borrowers, and the result is presented in Figure C.2b.

We use the reported numbers in [84] as a proxy for the debt-payment-to-income ratio of student loan borrowers for different family-income categories along with the estimated income distribution in Figure C.2b, and each of the proposed models to calculate the probabilities of default for each income category and the overall population of student loan borrowers. We then minimize the distance of these estimated probabilities from the empirical probabilities to determine the optimal values of α and β for each model.

[2]The Stafford loans are a part of the Direct Loan program, formerly known as the Federal Family Education Loan Program (FFELP).

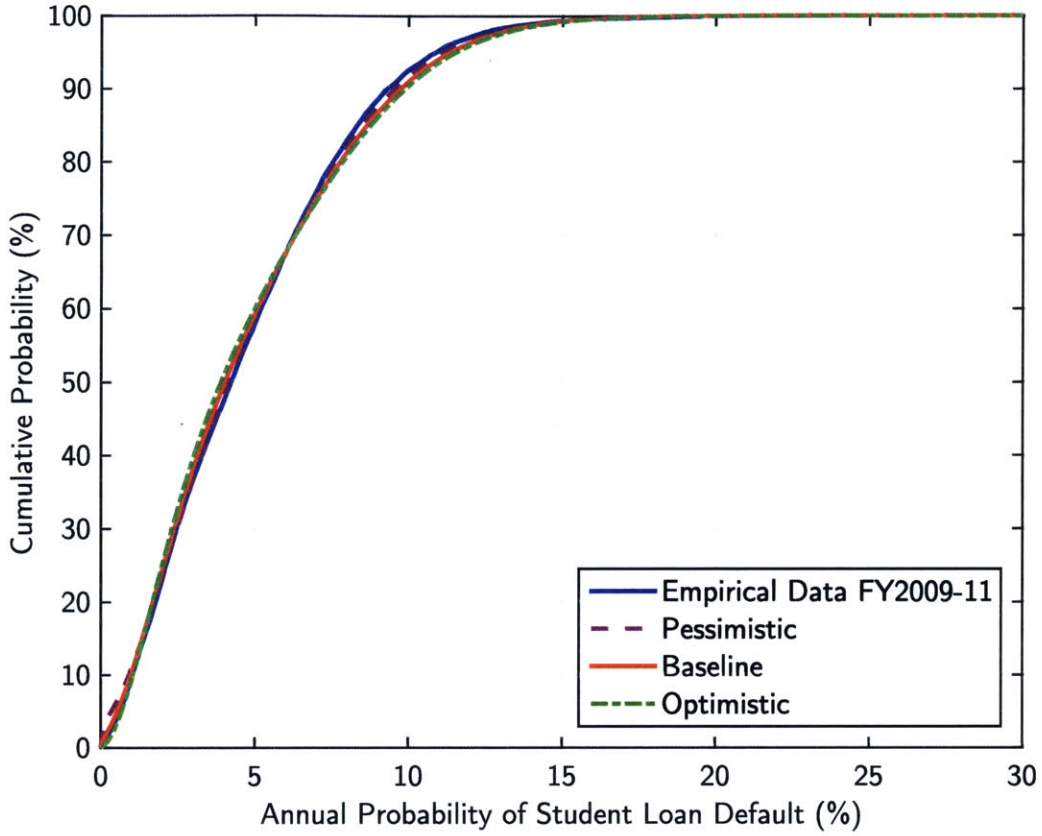

Figure C.3. The empirical distribution of federal student loan annual defaults for the fiscal years 2009–2011, and the resulting distributions from the models given by (C.1), (C.2), and (C.3) with their parameter values tabulated in Table C.1.

Table C.1. The optimal values of the two parameters, α and β, for the models given by (C.1), (C.2), and (C.3).

	Pessimistic	Baseline	Optimistic
α^*	0.28	1.24	2.02
β^*	0.11	1.70	0.48

These optimal values for each scenario are tabulated in Table C.1, and we present the resulting default probability distributions for each model using these optimal parameter values in Figure C.3.

■ C.3 Post-Treatment Survival Curve Estimation

In this section, we use the projected population numbers for the baby-boom cohort in the U.S., published by the U.S. Census Bureau [81], to derive a post-treatment survival curve for the HCV patients in the U.S. We focus on this specific cohort because an estimated 75%+ of the patients with HCV infection in the U.S. are baby-boomers. If the estimated population size at year t is given by $N(t)$, we define the survival curve of the population as the following:

$$S(t) \equiv \frac{N(t)}{N(2015)}, \quad t \geq 2015. \tag{C.4}$$

The population numbers, $N(t)$'s, are reported over years with 5-year increments, i.e., for $t = 2015, 2020, 2025,$, etc., and we interpolate the curve between every two adjacent samples that are 5-year apart. By inspecting the given survival rates for the selected years and their logarithms, we propose the following functional form:

$$\widehat{S}(t) = \exp\left(-\Lambda(t)\right), \quad t \geq 2015, \tag{C.5}$$

where $\exp(\cdot)$ is the exponential function, and the cumulative hazard function, $\Lambda(t)$, is given by

$$\Lambda(t + 2015) = -1 + \exp\left(at^2 + bt\right), \quad t \geq 0. \tag{C.6}$$

To avoid over-fitting when estimating the survival curve for the first 9 years (the time period required for our simulation), we use the survival numbers for $t = 2015, 2020, 2025$, 2030, and 2035; i.e., we consider the survival curve for 11 more years than required for our simulations. We then determine the values of parameters a and b in (C.6) such that they jointly minimize the error function, given by

$$E(a, b) = \sum_{k=0}^{4} \left(S(2015 + 5k) - \widehat{S}(2015 + 5k)\right)^2. \tag{C.7}$$

The minimizing values are $a^* = 5.0 \times 10^{-4}$ and $b^* = 57.5 \times 10^{-4}$, and the estimated survival curve is presented in Figure C.4.

To verify the accuracy of our estimated survival curve, we compare our estimate to the result obtained by using the 3-parameter Burr Type XII distribution [85] instead of the functional form given by (C.5) and (C.6). The Burr distribution is commonly used in survival analyses and encompasses twelve different distributions, including the

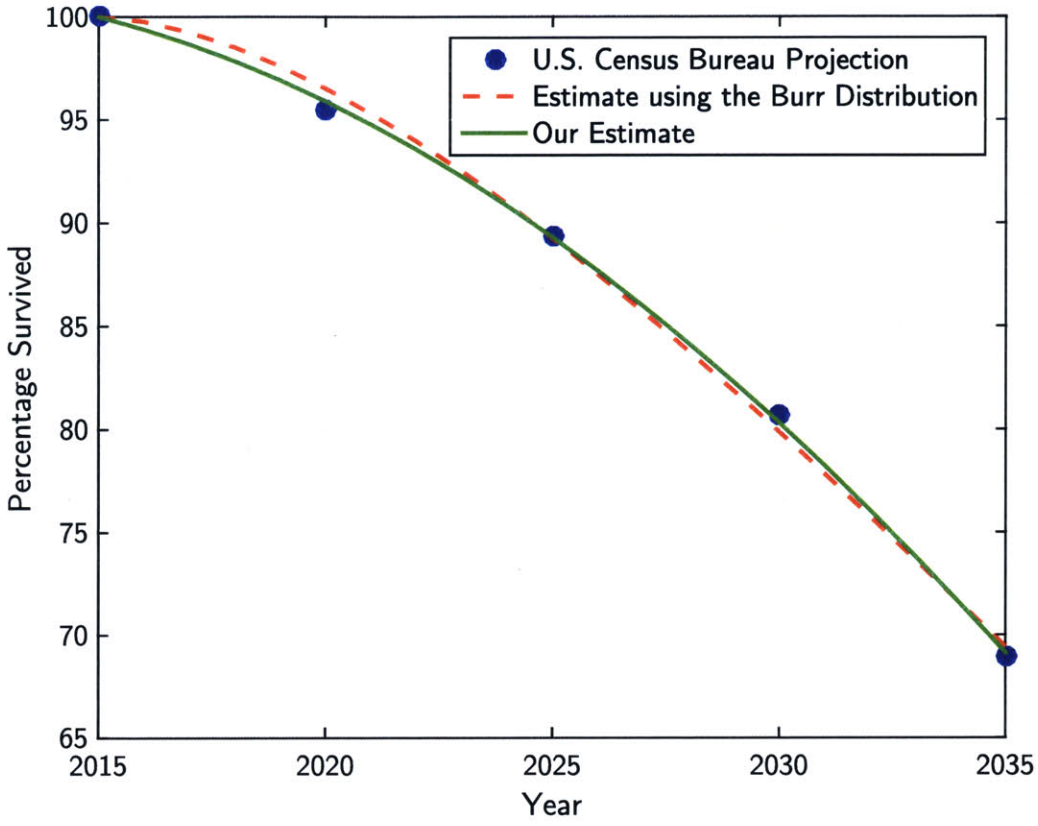

Figure C.4. Estimated Survival curves using the model given by (C.5) and (C.6), and using the Burr Type XII distribution. The projected population trend for the U.S. baby-boom cohort by the U.S. Census Bureau is also presented.

Weibull, logistic, and log-logistic distributions, as its special cases [85, 86]. The obtained results using this distribution are drawn in Figure C.4. The estimate error using the Burr distribution is higher than the estimate error using our proposed model in (C.5) and (C.6), and the results of the Burr distribution are a little more optimistic (i.e., yielding higher survival rates) over the first 9 years. However, as seen in Figure C.4, both estimates are pretty close.

Bibliography

[1] M. Hu and K. Schultz and J. Sheu and D. Tschopp. The innovation gap in pharmaceutical drug discovery & new models for R&D success. Available at `http://www.kellogg.northwestern.edu/biotech/faculty/articles/newrdmodel.pdf`. Accessed July 30, 2015, Mar. 2007.

[2] S. M. Paul, D. S. Mytelka, C. T. Dunwiddie, C. C. Persinger, B. H. Munos, S. R. Lindborg, and A. L. Schacht. How to improve R&D productivity: The pharmaceutical industry's grand challenge. *Nat Rev Drug Discov*, 9(3):203–214, Mar. 2010.

[3] F. Pammolli, L. Magazzini, and M. Riccaboni. The productivity crisis in pharmaceutical R&D. *Nat Rev Drug Discov*, 10(6):428–438, Jun. 2011.

[4] J. W. Scannell, A. Blanckley, H. Boldon, and B. Warrington. Diagnosing the decline in pharmaceutical R&D efficiency. *Nat Rev Drug Discov*, 11(3):191–200, Mar. 2012.

[5] U.S. Food and Drug Administration. Paving the way for personalized medicine: Fda's role in a new era of medical product development. Accessed July 28, 2015 at `http://www.fda.gov/downloads/ScienceResearch/SpecialTopics/PersonalizedMedicine/UCM372421.pdf`, Oct. 2013.

[6] B. Booth and R. Zemmel. Prospects for productivity. *Nat Rev Drug Discov*, 3(5): 451–456, May 2004.

[7] B. Dennis. Drugmakers find breakthroughs in medicine tailored to individuals' genetic makeups. The Washington Post. Available at `http://wpo.st/kjST0`. Accessed July 30, 2015, Jun. 2014.

[8] K. Polyak. Heterogeneity in breast cancer. *J Clin Invest*, 121(10):3786–3788, Oct. 2011.

[9] N. R. Bertos and M. Park. Breast cancer—one term, many entities? *J Clin Invest*, 121(10):3789–3796, Oct. 2011.

[10] A. Marusyk, V. Almendro, and K. Polyak. Intra-tumour heterogeneity: A looking glass for cancer? *Nat Rev Cancer*, 12(5):323–334, May 2012.

[11] National Venture Capital Association. National venture capital association yearbook 2014. Prepared by Thomson Reuters, New York, NY, 2014. Available at `http://www.nvca.org/index.php?option=com_docman&task=doc_download&gid=1054`. Accessed December 14, 2014.

[12] N. Nichols. Scientific management at Merck–an interview with CFO Judy Lewent. *Harvard Business Rev*, pages 89–99, Jan. 1994.

[13] J.-M. Fernandez, R. M. Stein, and A. W. Lo. Commercializing biomedical research through securitization techniques. *Nat Biotechnol*, 30(10):964–975, Oct. 2012.

[14] D. E. Fagnan, A. A. Gromatzky, R. M. Stein, J.-M. Fernandez, and A. W. Lo. Financing drug discovery for orphan diseases. *Drug Disc Today*, 19(5):533–538, May 2014.

[15] D. Thomas and C. Wessel. Venture funding of therapeutic innovation: A comprehensive look at a decade of venture funding of drug R&D. Biotechnology Industry Organization (BIO), Washington, DC. Accessed September 1, 2015 at `https://www.bio.org/sites/default/files/BIO-Whitepaper-FINAL.PDF`, Feb. 2015.

[16] H. Moses, D. H. M. Matheson, S. Cairns-Smith, B. P. George, C. Palisch, and E. R. Dorsey. The anatomy of medical research: US and international comparisons. *JAMA*, 313(2):174–189, Jan. 2015.

[17] B. Booth. The venture funding boom in biotech: A few things it's not. Forbes Magazine. Accessed July 30, 2015 at `http://onforb.es/1SDXQhU`, Jul. 2015.

[18] B. Huggett. Biotech's wellspring—a survey of the health of the private sector in 2014. *Nat Biotechnol*, 33(5):470–477, May 2015.

[19] Moody's Investors Service. Rating methodology: Moody's global approach to rating collateralized loan obligations, 2014.

[20] D. E. Fagnan, J.-M. Fernandez, A. W. Lo, and R. M. Stein. Can financial engineering cure cancer? *Amer Econ Rev*, 103(3):406–411, May 2013.

[21] D. E. Fagnan, N. N. Yang, J. C. McKew, and A. W. Lo. Financing translation: Analysis of the NCATS rare-diseases portfolio. *Sci Transl Med*, 7(276):276ps3, Feb. 2015.

[22] A. W. Lo and S. V. Naraharisetti. New financing methods in the biopharma industry: A case study of Royalty Pharma. *J Invest Manag*, 12(1):4–19, Jul. 2014.

[23] S. M. Forman, A. W. Lo, M. Shilling, and G. K. Sweeney. Funding translational medicine via public markets: The business development company. *J Invest Manag*, 13(4):1–24, Nov. 2015.

[24] F. S. David, S. T. Bobulsky, K. F. Schulz, and N. Patel. Creating value with financially adaptive clinical trials. *Nat Rev Drug Discov*, 14(8):523–524, Aug. 2015.

[25] K. F. Schulz, S. T. Bobulsky, F. S. David, N. Patel, and Z. Antonijevic. Drug development and the cost of capital. In Z. Antonijevic, editor, *Optimization of Pharmaceutical R&D Programs and Portfolios: Design and Investment Strategy.* Springer, New York, 1^{st} edition, 2015.

[26] S. J. Pocock. *Clinical Trials: A Practical Approach.* Wiley, New York, 1983.

[27] L. M. Friedman, C. D. Furberg, and D. L. DeMets. *Fundamentals of Clinical Trials: A Practical Approach.* Springer, New York, 4th edition, 2010.

[28] D. A. Berry. Bayesian clinical trials. *Nat Rev Drug Discov*, 5(1):27–36, Jan. 2006.

[29] U.S. Food and Drug Administration. Guidance for industry: Fast track drug development programs—designation, development, and application review. Accessed April 20, 2015 at `http://www.fda.gov/downloads/Drugs/Guidances/ucm079736.pdf`, Jan. 2006.

[30] U.S. Food and Drug Administration. Guidance for industry: Expedited programs for serious conditions— drugs and biologics. Accessed April 20, 2015 at `http://www.fda.gov/downloads/Drugs/GuidanceComplianceRegulatoryInformation/Guidances/UCM358301.pdf`, June 2013.

[31] M. Greener. Drug safety on trial. *EMBO Rep*, 6(3):202–204, Mar. 2005.

[32] J. K. Aronson. Drug withdrawals because of adverse effects. In J.K. Aronson, editor, *A worldwide yearly survey of new data and trends in adverse drug reactions and interactions*, volume 30 of *Side Effects of Drugs Annual*, pages xxxi–xxxv. Elsevier, 2008. doi: http://dx.doi.org/10.1016/S0378-6080(08)00064-0.

[33] R. McNaughton, G. Huet, and S. Shakir. An investigation into drug products withdrawn from the EU market between 2002 and 2011 for safety reasons and the evidence used to support the decision-making. *BMJ Open*, 4(1):e004221, Jan. 2014.

[34] ProCon.org. 35 FDA-approved prescription drugs later pulled from the market. Accessed March 18, 2015 at `http://prescriptiondrugs.procon.org/view.resource.php?resourceID=005528`, Jan. 2014.

[35] L. A. Lenert, D. R. Markowitz, and T. F. Blaschke. Primum non nocere? valuing of the risk of drug toxicity in therapeutic decision making. *Clin Pharmacol Ther*, 53(3):285–291, Mar. 1993.

[36] H.-G. Eichler, E. Abadie, J. M. Raine, and T. Salmonson. Safe drugs and the cost of good intentions. *New Engl J Med*, 360(14):1378–1380, Apr. 2009.

[37] H.-G. Eichler, B. Bloechl-Daum, D. Brasseur, and *et al.* The risks of risk aversion in drug regulation. *Nat Rev Drug Discov*, 12(12):907–916, Dec. 2013.

[38] F. J. Anscombe. Sequential medical trials. *J Am Stat Assoc*, 58(302):365–383, 1963.

[39] T. Colton. A model for selecting one of two medical treatments. *J Am Stat Assoc*, 58(302):388–400, 1963.

[40] D. A. Berry. Interim analysis in clinical trials: The role of likelihood principle. *Am Stat*, 41(2):117–122, May 1987.

[41] D. J. Spiegelhalter, L. S. Freedman, and M. K. B. Parmar. Bayesian approaches to randomized trials. *J R Stat Soc Ser A Stat Soc*, 157(3):357–416, 1994.

[42] M. H. DeGroot. *Optimal Statistical Decisions.* McGraw-Hill series in probability and statistics. McGraw-Hill, New York, 1970.

[43] C. J. L. Murray, J. Abraham, M. K. Ali, and *et al.* The state of U.S. health, 1990–2010: Burden of diseases, injuries, and risk factors. *JAMA*, 310(6):591–608, Jul. 2013.

[44] K. Y. Bilimoria, D. J. Bentrem, C. Y. Ko, and *et al.* Validation of the 6^{th} edition AJCC pancreatic cancer staging system. *Cancer*, 110(4):738–744, Aug. 2007.

[45] National Institute for Health and Care Excellence. Quality adjusted life years (QALYs) and severity of illness: Report 10. Accessed July 9, 2015 at `https://www.nice.org.uk/Media/Default/Get-involved/Citizens-Council/Reports/CCReport10QALYSeverity.pdf`, Feb. 2008.

[46] U.S. Food and Drug Administration. Draft PDUFA V implementation plan: Structured approach to benefit-risk assessment in drug regulatory decision-making. Accessed April 20, 2014 at `http://www.fda.gov/downloads/ForIndustry/UserFees/PrescriptionDrugUserFee/UCM329758.pdf`, Feb. 2013. Fiscal Years 2013-2017.

[47] U.S. Food and Drug Administration. Federal register notice. Accessed April 20, 2014 at `http://www.gpo.gov/fdsys/pkg/FR-2013-04-11/pdf/2013-08441.pdf`, Apr. 2013.

[48] U.S. Congress. Title 21, Code of Federal Regulations, part 312, subpart E: Drugs intended to treat life-threatening and severely-debilitating illnesses. Accessed April 20, 2014 at `http://www.gpo.gov/fdsys/pkg/CFR-1999-title21-vol5/pdf/CFR-1999-title21-vol5-part312-subpartE.pdf`, Apr. 1999.

[49] Center for Devices and Radiological Health of the U.S. Food and Drug Administration. Guidance for industry and FDA staff: Guidance for the use of Bayesian statistics in medical device clinical trials. Accessed March 14, 2015 at `http://www.fda.gov/downloads/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm071121.pdf`, Feb. 2010.

[50] D. A. Berry. Bayesian statistics and the efficiency and ethics of clinical trials. *Stat Sci*, 19(1):175–187, 2004.

[51] Y. Cheng, F. Su, and D. A. Berry. Choosing sample size for a clinical trial using decision analysis. *Biometrika*, 90(4):923–936, Dec. 2003.

[52] C. Jennison and B. W. Turnbull. *Group Sequential Methods with Applications to Clinical Trials.* CRC Press, 2010.

[53] B. Freedman. Equipoise and the ethics of clinical research. *New Engl J Med*, 317 (3):141–145, Jul. 1987.

[54] A. D. Barker, C. C. Sigman, G. J. Kelloff, N. M. Hylton, D. A. Berry, and L. J. Esserman. I-SPY 2: An adaptive breast cancer trial design in the setting of neoadjuvant chemotherapy. *Clin Pharmacol Ther*, 86(1):97–100, May 2009.

[55] J. A. Salomon, T. Vos, D. R. Hogan, and *et al.* Common values in assessing health outcomes from disease and injury: Disability weights measurement study for the Global Burden of Disease Study 2010. *Lancet*, 380(9859):2129–2143, Dec. 2012.

[56] O. B. Ahmad and C. Boschi-Pinto and A. D. Lopez and *et al.* Age standardization of rates: A new WHO standard. GPE discussion paper series: No 31, Accessed July 20, 2014 at `http://www.who.int/healthinfo/paper31.pdf`, 2001.

[57] D. A. Berry. The Brave New World of clinical cancer research: Adaptive biomarker-driven trials integrating clinical practice with clinical research. *Mol Oncol*, 9(5): 951–959, Mar. 2015.

[58] U.S. Congress. Food and Drug Administration Safety and Innovation Act, public law 112-144. Accessed April 20, 2014 at `http://www.gpo.gov/fdsys/pkg/PLAW-112publ144/pdf/PLAW-112publ144.pdf`, Jul. 2012.

[59] H.-G. Eichler, E. Abadie, A. Breckenridge, and *et al.* Bridging the efficacy-effectiveness gap: A regulator's perspective on addressing variability of drug response. *Nat Rev Drug Discov*, 10(7):495–506, Jul. 2011.

[60] H. P. Selker, K. A. Oye, H.-G. Eichler, and *et al.* A proposal for integrated efficacy-to-effectiveness (E2E) clinical trials. *Clin Pharmacol Ther*, 59(2):147–153, Feb. 2014.

[61] H.-G. Eichler, F. Pignatti, B. Flamion, H. Leufkens, and A. Breckenridge. Balancing early market access to new drugs with the need for benefit/risk data: A mounting dilemma. *Nat Rev Drug Discov*, 7(10):818–826, Oct. 2008.

[62] European Medicines Agency. Road map to 2015: The European medicines agency's contribution to science, medicines and health. Accessed March 15, 2015 at `http://www.ema.europa.eu/docs/en_GB/document_library/Report/2011/01/WC500101373.pdf`, Dec. 2010.

[63] H.-G. Eichler, K. Oye, L. G. Baird, and *et al.* Adaptive licensing: Taking the next step in the evolution of drug approval. *Clin Pharmacol Ther*, 91(3):426–437, Mar. 2012.

[64] European Medicines Agency. Adaptive pathways to patients: Report on the initial experience of the pilot project. Technical Report EMA/758619/2014, Accessed July 20, 2015 at `http://www.ema.europa.eu/docs/en_GB/document_library/Report/2014/12/WC500179560.pdf`, European Medicines Agency, Dec. 2014.

[65] L. G. Baird, M. R. Trusheim, H.-G. Eichler, E. R. Berndt, and G. Hirsch. Comparison of stakeholder metrics for traditional and adaptive development and licensing approaches to drug development. *Ther Innov Regul Sci*, 47(4):474–483, Jul. 2013.

[66] E. Lawitz, A. Mangia, D. Wyles, and *et al.* Sofosbuvir for previously untreated chronic hepatitis C infection. *New Engl J Med*, 368(20):1878–1887, 2013.

[67] R. L. Kosowski and S. Neftci. *Principles of Financial Engineering.* Academic Press, San Diego, CA, 3^{rd} edition, 2014.

[68] T. A. Brennan and J. M. Wilson. The special case of gene therapy pricing. *Nat Biotechnol*, 32(9):874–876, Sep. 2014.

[69] S. A. Ross, R. Westerfield, and J. D. Bradford. *Fundamentals of Corporate Finance.* McGraw-Hill Irwin, Boston, MA, 2012.

[70] R. A. Brealey, S. C. Myers, and A. J. Marcus. *Fundamentals of Corporate Finance.* McGraw-Hill/Irwin, New York, 2012.

[71] H. Markowitz. Portfolio selection. *J Finance*, 7(1):77–91, Mar. 1952.

[72] H. M. Markowitz. *Portfolio Selection: Efficient Diversification of Investments.* B. Balckwell, Cambridge, MA, 2^{nd} edition, 1991.

[73] B. P. Lancaster, G. M. Schultz, and F. J. Fabozzi. *Structured Products and Related Credit Derivatives: A Comprehensive Guide for Investors*, volume 151 of *The Frank J. Fabozzi Series.* John Wiley & Sons, Hoboken, NJ, 2008.

[74] J. Bricker, L. J. Dettling, A. Henriques, and *et al.* Changes in U.S. family finances from 2010 to 2013: Evidence from the survey of consumer finances. *Federal Reserve Bulletin*, 100(4):1–41, 2014.

[75] A. C. Moorman, S. C. Gordon, L. B. Rupp, and *et al.* Baseline characteristics and mortality among people in care for chronic viral hepatitis: The chronic hepatitis cohort study. *Clin Infect Dis*, 56(1):40–50, 2013.

[76] K. Yamasaki, M. Tomohiro, Y. Nagao, and *et al.* Effects and outcomes of interferon treatment in Japanese hepatitis C patients. *BMC Gastroenterol*, 12(139), 2012.

[77] B. Simmons, J. Saleem, K. Heath, G. S. Cooke, and A. Hill. Long-term treatment outcomes of patients infected with hepatitis C virus: A systematic review and meta-analysis of the survival benefit of achieving a sustained virological response. *Clin Infect Dis*, civ396:1–11, 2015.

[78] J. Smith-Palmer, K. Cerri, and W. Valentine. Achieving sustained virologic response in hepatitis C: A systematic review of the clinical, economic and quality of life benefits. *BMC Infect Dis*, 15(19), Jan. 2015.

[79] M.-H. Lee, H.-I. Yang, S.-N. Lu, and *et al.* Chronic hepatitis C virus infection increases mortality from hepatic and extrahepatic diseases: A community-based long-term prospective study. *J Infect Dis*, 206(4):469–477, 2012.

[80] A. J. van der Meer and B. E. Hansen and J. J. Feld and *et al.* Comparison of the overall survival between patients with HCV-induced advanced hepatic fibrosis and the general population. 64^{th} Annual Meeting of the American Association for the Study of Liver, Washington, DC, 2013.

[81] S. L. Colby and J. M. Ortman. The baby boom cohort in the United States: 2012 to 2060. Current Population Reports, P25 – 1141. U.S. Census Bureau, Washington, DC, May 2014.

[82] A. J. van der Meer, B. J. Veldt, J. J. Feld, and *et al.* Association between sustained virological response and all-cause mortality among patients with chronic hepatitis C and advanced hepatic fibrosis. *JAMA*, 308(24):2584–2593, 2012.

[83] National Association of Student Financial Aid Administrators. National student aid profile: Overview of 2012 federal programs. NASFAA Report, Washington, DC, 2012.

[84] R. Fry. A record one-in-five households now owe student loan debt. Pew Social & Demographic Trends, Washington, DC, 2012.

[85] I. W. Burr. Cumulative frequency functions. *Ann Math Stat*, 13(2):215–232, 1942.

[86] R. N. Rodriguez. A guide on the Burr type XII distributions. *Biometrika*, 64(1): 120–134, 1977.